arXiv:1702.07810v1 [cs.GT] 25 Feb 2017

A Decomposition of Forecast Error in Prediction Markets

Miroslav Dudík[1], Sébastien Lahaie[1], Ryan Rogers[2], and Jennifer Wortman Vaughan[1]

[1]Microsoft Research
[2]University of Pennsylvania

February 28, 2017

Abstract

We introduce and analyze sources of error in prediction market forecasts in order to characterize and bound the difference between a security's price and its ground truth value. We consider cost-function-based prediction markets in which an automated market maker adjusts security prices according to the history of trade. We decompose the forecasting error into four components: *sampling error*, occurring because traders only possess noisy estimates of ground truth; *risk-aversion effect*, arising because traders reveal beliefs only through self-interested trade; *market-maker bias*, resulting from the use of a particular market maker (i.e., cost function) to facilitate trade; and finally, *convergence error*, arising because, at any point in time, market prices may still be in flux. Our goal is to understand the tradeoffs between these error components, and how they are influenced by design decisions such as the functional form of the cost function and the amount of liquidity in the market. We specifically consider a model in which traders have exponential utility and exponential-family beliefs drawn with an independent noise relative to ground truth. In this setting, sampling error and risk-aversion effect vanish as the number of traders grows, but there is a tradeoff between the other two components: decreasing the market maker's liquidity results in smaller market-maker bias, but may also slow down convergence. We provide both upper and lower bounds on market-maker bias and convergence error, and demonstrate via numerical simulations that these bounds are tight. Our results yield new insights into the question of how to set the market's liquidity parameter, and into the extent to which markets that enforce coherent prices across securities produce better predictions than markets that price securities independently.

Contents

1 Introduction **4**

2 Preliminaries **6**
2.1 Cost-function-based Market Makers 6
2.2 Basic Components of the Model 7

3 A Decomposition of Error **8**

4 The Exponential Trader Model **9**
4.1 Trader Model 9
4.2 Potential Functions and Uniqueness of Equilibrium Prices 10

5 Sampling Error and the Risk-Aversion Effect **11**

6 Market-maker Bias **12**
6.1 Smoothness and Strict Convexity Assumptions 12
6.2 Bounding the Market-Maker Bias 13
6.3 Tighter Analysis for Convex$^+$ Cost Functions 14

7 Convergence Error **15**
7.1 Trader Dynamics 15
7.2 Global Convergence Rate 16
7.3 Local Convergence Rate of Block-coordinate Descent 17
7.4 Local Convergence Rate of the Market 18

8 Numerical experiments **19**

References **23**

A Mathematical background **24**
A.1 Vectors, matrices, intervals 24
A.2 Convex analysis 25
A.3 Convex$^+$ functions 26
A.4 Lipschitz Gradients and Strong Convexity 31

B Proofs from Section 4 **32**
B.1 Proof of Theorem 4.1 32
B.2 Proof of Theorem 4.2 33

C Proofs from Section 5 **34**
C.1 Proof of Theorem 5.1 34
C.2 Proof of Theorem 5.2 35

D Proofs and Additional Results from Section 6 **35**
D.1 Proof of Theorem 6.5 . . . 35
D.2 A Remark on Partial and Incoherent Beliefs . . . 35
D.3 Proof of Theorem 6.6 . . . 36

E Proofs and Additional Results for Section 7 **37**
E.1 Relationship between the Suboptimality of Potential and the Convergence Error . . . 37
E.2 Proof of Theorem 7.2 . . . 39
E.3 Local Convergence Rate of Block-Coordinate Descent . . . 39
E.4 Supporting Propositions for Theorem 7.5 . . . 43
E.5 Proof of Theorem 7.5 . . . 45
E.6 Local Convergence of the Market . . . 46
E.7 Tighter Analysis of the All-securities Dynamics . . . 50
E.8 Proof of Theorem 7.6 . . . 56

F Additional Results from Numerical Experiments **56**
F.1 Tightness of the Empirical Strong Convexity for Single-Peaked Beliefs . . . 56
F.2 Numerical Experiments for Single-Security Dynamics . . . 56

1 Introduction

A prediction market is a marketplace in which participants can trade securities with payoffs that depend on the outcomes of future events Wolfers and Zitzewitz [2004]. Consider the simple setting in which we are interested in predicting the outcome of a political election, that is, whether the incumbent or challenger will win. A prediction market might issue a security that pays out \$1 per share if the incumbent wins, and \$0 per share otherwise. As the market price p lies between 0 and 1, it can be construed as the event probability. If a trader believes that the chances of the incumbent winning are greater than p, she will buy shares with the expectation of making a profit, pushing the security price closer to her belief. By this process, the market aggregates traders' information into a consensus forecast, represented by the market price. With sufficient activity, prediction markets are competitive with alternative prediction methods such as polls Berg et al. [2008], but while there is a mature literature on sources of error and bias in polls, the impact of prediction market structure on forecast accuracy is still an active area of research Rothschild [2009].

We consider prediction markets in which all trades occur through a centralized entity known as a *market maker*. Under this market structure, security prices are dictated by a fixed *cost function* and the current number of outstanding shares Chen and Pennock [2007]. The basic conditions that a cost function should satisfy, in order to bound the market maker's loss and avoid arbitrage, are now well-understood, chief among them being convexity Abernethy et al. [2013]. Nonetheless, the class of allowable cost functions remains broad, and the literature so far provides little formal guidance on the specific form of cost function to use in order to achieve good forecast accuracy. Even once a cost function is selected, there is still a free *liquidity parameter* which controls price responsiveness to trade, and therefore impacts how prices incorporate information. In practice, the impact of the liquidity parameter is difficult to quantify a priori, and implementations typically resort to calibrations based on market simulations [e.g, Dudík et al., 2013, Slamka et al., 2013].

This paper provides a framework to quantify the impact of market structure, including the liquidity level used and the form of cost function more generally, on forecast accuracy. We introduce a decomposition of forecast error in cost-function-based prediction markets, in analogy with the bias-variance decomposition familiar from statistics. Our decomposition consists of four components. First, there is the *sampling error* that stems from the fact that the market consists of a finite population of traders, each holding a noisy estimate of the ground truth. Second, this error may be compounded by a *risk-aversion effect*, due to the fact that varying risk attitudes among the traders can unduly bias market prices towards certain private signals. Third, there is a *market-maker bias* which stems from the use of a cost function to provide liquidity and induce trade. Intuitively, a market with high liquidity (low price responsiveness) may not adequately incorporate the traders' information into the market prices; the actual form of the cost function also impacts this bias in non-obvious ways. Fourth, there is *convergence error* due to the fact that, at any particular point in time, the market prices may not have fully converged to their equilibrium point.

The central contribution of this paper is a theoretical characterization of each component of this decomposition, in terms of the dependence on the form of the cost function and level of liquidity. Taken together, our results quantify the tradeoffs inherent in various choices of liquidity levels and cost functions. We consider a tractable model of agent behavior, originally studied by Abernethy et al. [2014], in which traders have exponential utility and beliefs drawn from an exponential family. Under this model it is possible to characterize the market's equilibrium prices in terms of traders' belief and risk aversion parameters, and therefore to quantify the discrepancy between current market prices and ground truth. Our analysis focuses most heavily on market-maker bias and

convergence error, because these are the components that depend on market structure.

In Section 2, we introduce relevant background on cost-function-based market makers. In Section 3, we provide a formal description of our error decomposition and relate the sources of error to certain notions of market equilibrium prices, which represent the consensus forecasts of a market with and without a market maker. Section 4 describes our model of traders, including exponential utility and exponential-family beliefs, and characterizes equilibrium prices with such traders. The remaining sections contain our novel results.

In Section 5, we bound the sampling error and risk-aversion effect. We confirm that as long as traders have bounded risk aversion, the discrepancy between market-clearing prices and ground truth goes to zero as the population of traders grows large. The bound increases with the noise in the traders' signals of ground truth, and with the spread in their risk aversion. These error components serve as a reminder that any market can do little to ensure good forecasts if the population of traders is too small, ill-informed, or too risk averse.

In Section 6, we analyze the market-maker bias. We provide both *global* and *local* bounds to characterize the dependence on liquidity and the cost function. A local bound depends on the location of the market equilibrium prices, and therefore trader beliefs, whereas a global bound holds independent of the realized beliefs. We prove a global bound on market-maker bias that decreases linearly in the liquidity level used in the market. This bound is too weak to distinguish between cost functions, however, so we prove a local characterization that depends on the shape (more formally, the Hessian) of the cost function at market equilibrium prices. We use our characterization to compare the bias of two common cost functions: independent markets (IND), under which security prices vary independently, and the logarithmic market scoring rule (LMSR), which enforces logical relationships between security prices. We find that under certain conditions, IND has up to twice the market-maker bias of LMSR at the same liquidity level.

In Section 7, we turn to the convergence error component. We consider two types of trader dynamics: one in which agents may trade all securities at each round, and another in which traders are limited to consider only a single security at each round. Both can be analyzed by viewing the market as a convex optimization algorithm performing block-coordinate descent. Prior work by Frongillo and Reid [2015] shows a global convergence bound of $O(1/t)$ after t rounds of trade, while conjecturing a stronger local result. We show that a sharper local bound of $O(\gamma^t)$ indeed holds with $\gamma = 1 - \Omega(b)$, where b is the liquidity parameter. Along the way we extend the results of Nesterov [2012] and prove a new local convergence bound for block-coordinate descent, which may be of independent interest. For all-securities dynamics, we give a matching lower bound on γ and thus fully characterize the convergence rate. Again, our local characterization allows us to compare the convergence rate of different cost functions, beyond just the effect of liquidity. We show tightness of our local bounds empirically.

In Section 8, we conduct extensive numerical simulations to corroborate our bounds on market-maker bias and convergence error. We find that our local characterization of bias is tight, which justifies its use in comparing cost functions. We also find that the local regime occurs early on, which justifies the use of the local analysis in comparing convergence error between cost functions. Our theoretical framework therefore provides a meaningful way to compare the detailed quantitative error tradeoffs inherent in different choices of cost functions and liquidity levels.

2 Preliminaries

We first review the cost-function-based market making framework [Chen and Pennock, 2007, Abernethy et al., 2013] and describe the basic components of our trader model.

Here and throughout the paper, vectors are denoted by small boldface italics (e.g., $\boldsymbol{s}$, $\boldsymbol{\mu}$), while matrices are denoted by large italics (e.g., A, P). For a pair of symmetric square matrices A and B we write $A \preceq B$ iff $\boldsymbol{u}^\intercal A \boldsymbol{u} \leq \boldsymbol{u}^\intercal B \boldsymbol{u}$ for all $\boldsymbol{u}$. We use the notation $[N]$ to denote the set $\{1, \ldots, N\}$. Given a convex function $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$, its *effective domain*, denoted $\operatorname{dom} f$, is the set of points where it is finite. Whenever its domain is non-empty, its *conjugate* $f^* : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ is defined by $f^*(\boldsymbol{v}) := \sup_{\boldsymbol{u} \in \mathbb{R}^d} [\boldsymbol{v}^\intercal \boldsymbol{u} - f(\boldsymbol{u})]$. Relevant mathematical background is introduced along the way as it is needed, but we refer the reader to Appendix A for a centralized mathematical reference.

2.1 Cost-function-based Market Makers

We study cost-function-based prediction markets in the general setting formalized by Abernethy et al. [2013]. Let Ω be a finite set of mutually exclusive and exhaustive states of the world. An automated market maker offers a set of K securities. The values of these securities are defined by an arbitrary payoff function $\boldsymbol{\phi} : \Omega \to \mathbb{R}^K$, with the kth security worth $\phi_k(\omega)$ in the event that the true state of the world is $\omega \in \Omega$.

Traders arrive in the market sequentially and trade with the market maker, who is always willing to buy or sell the K securities at some current market price that depends on the history of trade. The market price is fully determined by a convex potential function C called the *cost function*. In particular, if the market maker has previously sold $s_k \in \mathbb{R}$ shares of each security k and a trader would like to purchase a bundle consisting of $\delta_k \in \mathbb{R}$ shares of each, the trader is charged $C(\boldsymbol{s} + \boldsymbol{\delta}) - C(\boldsymbol{s})$. The *instantaneous price* of security k is then $\partial C(\boldsymbol{s}) / \partial s_k$. Note that negative values of δ_k are allowed and correspond to the trader (short) selling security k.

Let $\mathcal{M} := \operatorname{conv}\{\boldsymbol{\phi}(\omega) : \omega \in \Omega\}$ be the convex hull of the set of payoff vectors. It is exactly the set of expectations $\mathbb{E}[\boldsymbol{\phi}(\omega)]$ across all possible probability distributions over Ω, which we call *beliefs*. We refer to elements of $\mathcal{M}$ as *coherent prices*. Abernethy et al. [2013] characterize the conditions that a cost function must satisfy in order to guarantee important properties such as bounded loss for the market maker and no possibility of arbitrage. Here, we assume only that $C : \mathbb{R}^K \to \mathbb{R}$ is convex, differentiable and that $\mathcal{M} \subseteq \operatorname{dom} C^*$, which corresponds to the bounded loss property. However, parts of our analysis make use of stronger assumptions on C.

Throughout this paper, we repeatedly return to several running examples of cost functions. The most commonly used cost function is Hanson's logarithmic market scoring rule (LMSR) [Hanson, 2003] which is defined for *complete markets* with $K = |\Omega|$.

Example 2.1 (Logarithmic Market Scoring Rule: LMSR)**.** Consider a complete market in which there is a single security for each outcome worth \$1 if that outcome occurs and \$0 otherwise; more formally, $\Omega = [K]$ and $\phi_k(\omega) = \mathbf{1}\{k = \omega\}$ for all k. The LMSR cost function and instantaneous security prices are given by

$$C(\boldsymbol{s}) = \log\left(\textstyle\sum_{k=1}^K e^{s_k}\right) \qquad \text{and} \qquad \frac{\partial C(\boldsymbol{s})}{\partial s_k} = \frac{e^{s_k}}{\sum_{\ell=1}^K e^{s_\ell}}, \ \forall k \in [K]. \tag{1}$$

Its conjugate is the entropy function, $C^*(\boldsymbol{\mu}) = \sum_k \mu_k \log \mu_k + \mathbb{I}\{\boldsymbol{\mu} \in \Delta_K\}$, where Δ_K is the simplex in $\mathbb{R}^K$ and $\mathbb{I}\{\cdot\}$ is the convex indicator, equal to zero if its argument is true and infinity if false. Thus, in this case $\mathcal{M} = \Delta_K = \operatorname{dom} C^*$.

Notice that the LMSR security prices are coherent because they always sum to one (i.e., they correspond to a probability distribution over the K outcomes). This prevents arbitrage opportunities for traders. The second running example of a cost function that we will use throughout this paper does not have this property. This cost function is also for a complete market with the same $K = |\Omega|$ securities. The costs it produces are equivalent to the costs that would be produced using separate binary LMSR cost functions to independently price each security. As a result, the price of security k does not change as other securities are bought and sold.

Example 2.2 (Sum of Independent LMSRs: IND)**.** Let $\Omega = [K]$ and $\phi_k(\omega) = \mathbf{1}\{k = \omega\}$ for all k. The cost function and instantaneous security prices for the *sum of independent LMSRs* are given by

$$C(\boldsymbol{s}) = \textstyle\sum_{k=1}^{K} \log\left(1 + e^{s_k}\right) \qquad \text{and} \qquad \dfrac{\partial C(\boldsymbol{s})}{\partial s_k} = \dfrac{e^{s_k}}{1 + e^{s_k}},\ \forall k \in [K]. \tag{2}$$

Its conjugate is the sum of binary entropies, $C^*(\boldsymbol{\mu}) = \sum_k [\mu_k \log \mu_k + (1 - \mu_k)\log(1 - \mu_k)] + \mathbb{I}\{\boldsymbol{\mu} \in [0, 1]^K\}$. In this case, $\mathcal{M} = \Delta_K$ while $\operatorname{dom} C^* = [0, 1]^K$.

When choosing a cost function, one important consideration is the *liquidity* of the resulting market, that is, how quickly prices change in response to trades. Any cost function C can be viewed as a member of a parametric family of cost functions of the form

$$C_b(\boldsymbol{s}) \coloneqq bC(\boldsymbol{s}/b)$$

for all $b > 0$. With larger values of b, larger trades are required to move market prices by some fixed amount, and the worst-case loss of the market maker is generally larger. Conversely, with smaller values of b, smaller purchases can result in big changes to the market price.

2.2 Basic Components of the Model

In order to discuss the aggregation of information and the error of market predictions, it is necessary to assume a model of trader behavior and to specify how traders' information relates to some notion of ground truth. To this end, in our analysis of error we assume that there exists an unknown true probability distribution $\boldsymbol{p}^{\text{true}} \in \Delta^{|\Omega|}$ over the outcome set Ω. The true expected payoffs of the K market securities are then given by a vector $\boldsymbol{\mu}^{\text{true}} \coloneqq \mathbb{E}_{\omega \sim \boldsymbol{p}^{\text{true}}}[\boldsymbol{\phi}(\omega)]$.

We assume that there are N traders trading in the market, and that each trader $i \in [N]$ has a private belief $\tilde{\boldsymbol{p}}_i$ over outcomes. The expected security payoffs under trader i's beliefs are given by the vector $\tilde{\boldsymbol{\mu}}_i \coloneqq \mathbb{E}_{\omega \sim \tilde{\boldsymbol{p}}_i}[\boldsymbol{\phi}(\omega)]$. For now we make no assumptions about the way in which $\tilde{\boldsymbol{p}}_i$ relates to the ground truth $\boldsymbol{p}^{\text{true}}$, though we will do so later when discussing sampling error.

We additionally assume that each trader i has a *utility function* $u_i : \mathbb{R} \to \mathbb{R}$ for wealth and would like to maximize expected utility subject to his beliefs. For now we assume that u_i is differentiable and concave, meaning that each trader is risk averse, though later we focus on exponential utility. The expected utility of trader i with a security bundle $\boldsymbol{r}_i \in \mathbb{R}^K$ and cash c_i is given by

$$U_i(\boldsymbol{r}_i, c_i) \coloneqq \mathbb{E}_{\omega \sim \tilde{\boldsymbol{p}}_i}\left[u_i\big(c_i + \boldsymbol{\phi}(\omega) \cdot \boldsymbol{r}_i\big)\right].$$

For convenience, we assume that each trader begins with zero cash endowment. This assumption is without loss of generality because we could incorporate any initial cash holdings into u_i.

3 A Decomposition of Error

We are interested in bounding the error of the market's predictions, that is, the difference between the market price of each security and the security's true expected value. As a first step, we show how this error can be decomposed into several components which we analyze in subsequent sections.

First, the *sampling error* arises because traders do not have perfect information but only noisy observations of ground truth. Even if each trader were to honestly reveal his information and this information were used to optimally infer each security's value (e.g., with a maximum likelihood estimate), there would still be error from noise. Second, there is what might be called a *risk-aversion effect* since traders do not report their beliefs truthfully but rather reveal them indirectly through self-interested trades. The magnitude of this error depends on the form of traders' utility functions.

Third, there is a *market-maker bias* term which arises because of the use of a cost-function-based market maker to induce trade. Intuitively, if a market has high liquidity, meaning that prices change only very slowly as securities are bought and sold, then risk averse traders may not be willing to invest the cash that would be needed to move the market prices close to their beliefs. The functional form of the cost function C may also affect bias in less obvious ways.

Finally, there is a *convergence error*, due to the fact that at any particular point in time the market prices may not have fully converged. The convergence error also depends on C and the liquidity b. For example, if liquidity is too low, traders can move the market prices close to their respective beliefs by making only small purchases, causing market prices to fluctuate as traders with different beliefs enter the market. As the traders' willingness to use their cash decreases, they eventually grow willing to accept prices further from their beliefs, allowing prices to converge. However, it takes many trades to reach this point, since each trade uses up only a small amounts of cash.

To define these error terms formally, we define two equilibrium notions. The *market-clearing equilibrium* is a combination of allocations $(\bar{\boldsymbol{r}}_i, \bar{c}_i)$ of security bundles and cash to each trader i and security prices $\bar{\boldsymbol{\mu}}$ such that, given her allocation, no trader would want to buy or sell any bundle of securities at the given prices.

Definition 3.1 (Market-clearing equilibrium)**.** A triple $(\bar{\boldsymbol{r}}, \bar{\boldsymbol{c}}, \bar{\boldsymbol{\mu}})$ is a market-clearing equilibrium if $\sum_{i=1}^{N} \bar{\boldsymbol{r}}_i = \mathbf{0}$, $\sum_{i=1}^{N} \bar{c}_i = 0$, and for all $i \in [N]$, $\mathbf{0} \in \operatorname{argmax}_{\boldsymbol{\delta} \in \mathbb{R}^K} U_i(\bar{\boldsymbol{r}}_i + \boldsymbol{\delta}, \bar{c}_i - \boldsymbol{\delta} \cdot \bar{\boldsymbol{\mu}})$. We call $\bar{\boldsymbol{\mu}}$ *market-clearing prices* if there exist some $\bar{\boldsymbol{r}}$ and $\bar{\boldsymbol{c}}$ such that $(\bar{\boldsymbol{r}}, \bar{\boldsymbol{c}}, \bar{\boldsymbol{\mu}})$ is a market-clearing equilibrium.

The requirements that $\sum_{i=1}^{N} \bar{\boldsymbol{r}}_i = \mathbf{0}$ and $\sum_{i=1}^{N} \bar{c}_i = 0$ guarantee that no additional securities or cash have been created; in other words, there exists some set of trades among traders that would lead to this allocation, although the definition itself says nothing about how the equilibrium is reached. In general market-clearing prices $\bar{\boldsymbol{\mu}}$ are not unique, but they will be unique for the specific utility functions that we study later in this paper.

Since we must rely on a market maker to orchestrate trade, the markets we consider generally do not reach the market-clearing equilibrium. Instead, we introduce the notion of a *market-maker equilibrium*. This equilibrium is again a combination of allocations $(\boldsymbol{r}_i^\star, c_i^\star)$ of security bundles and cash to each trader i and a set of security prices $\boldsymbol{\mu}^\star$ such that no trader would want to trade at these prices given his allocation. The difference is that we now require $\boldsymbol{r}^\star$ and $\boldsymbol{c}^\star$ to be reachable via some sequence of trade with the market maker instead of via trade among only the traders. The prices $\boldsymbol{\mu}^\star$ are precisely the prices the market maker would charge after this sequence of trade.

Definition 3.2 (Market-maker equilibrium)**.** A triple $(\boldsymbol{r}^\star, \boldsymbol{c}^\star, \boldsymbol{\mu}^\star)$ is a market-maker equilibrium for cost function C with liquidity parameter b if, for the market state $\boldsymbol{s}^\star = \sum_{i=1}^N \boldsymbol{r}_i^\star$,

$$\sum_{i=1}^N c_i^\star + C_b(\boldsymbol{s}^\star) - C_b(\mathbf{0}) = 0, \quad \boldsymbol{\mu}^\star = \nabla C_b(\boldsymbol{s}^\star), \quad \mathbf{0} \in \operatorname*{argmax}_{\boldsymbol{\delta} \in \mathbb{R}^K} U_i\big(\boldsymbol{r}_i^\star + \boldsymbol{\delta},\, c_i^\star - C_b(\boldsymbol{s}^\star + \boldsymbol{\delta}) + C_b(\boldsymbol{s}^\star)\big) \quad \forall i \in [N].$$

We call the prices $\boldsymbol{\mu}^\star$ *market-maker equilibrium prices* for C_b if there exist some $\boldsymbol{r}^\star$ and $\boldsymbol{c}^\star$ such that $(\boldsymbol{r}^\star, \boldsymbol{c}^\star, \boldsymbol{\mu}^\star)$ is a market-maker equilibrium for C_b. We sometimes write $\boldsymbol{\mu}^\star(b; C)$ to emphasize the dependence on the cost function family C and liquidity parameter b.

Like the market-clearing prices, the market-maker equilibrium prices $\boldsymbol{\mu}^\star(b; C)$ are not unique in general, but will be unique for the specific utility functions that we study in this paper.

Using these notions of equilibrium, we can formalize the error decomposition described above. The combination of sampling error and risk-aversion bias can be expressed as the difference between the true security values and the market-clearing equilibrium prices. The bias is the difference between the market-clearing equilibrium prices and the market-maker equilibrium prices. Finally, the convergence error is the difference between the market-maker equilibrium prices and the market prices $\boldsymbol{\mu}^t(b; C)$ at a particular round t. Putting this together, we have that

$$\boldsymbol{\mu}^{\text{true}} - \boldsymbol{\mu}^t = \underbrace{\boldsymbol{\mu}^{\text{true}} - \bar{\boldsymbol{\mu}}}_{\text{Sampling+Risk-aversion}} + \underbrace{\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}^\star(b; C)}_{\text{Bias}} + \underbrace{\boldsymbol{\mu}^\star(b; C) - \boldsymbol{\mu}^t(b; C)}_{\text{Convergence}}. \tag{3}$$

Since our primary goal is to understand the impact of the cost function C and liquidity parameter b on error, we devote most of our attention to analyzing the market-maker bias and convergence error, though we also briefly explore the sampling error and risk-aversion bias.

4 The Exponential Trader Model

We showed that the error in market predictions can be decomposed into several terms. For the remainder of the paper, in order to quantify these sources of error and how they are affected by the choice of cost function C, we impose further assumptions on the traders in our model. In particular, following Abernethy et al. [2014], we assume that traders have exponential utility functions and exponential-family beliefs. Under these assumptions, both the market-clearing prices and market-maker equilibrium prices are unique (though the equilibrium allocations might not be) and can be expressed cleanly in terms of a potential function, yielding a tractable analysis.

4.1 Trader Model

An exponential family is a parameterized family of probability distributions of the form $p(\omega; \boldsymbol{\theta}) = e^{\boldsymbol{\phi}(\omega)\cdot\boldsymbol{\theta} - T(\boldsymbol{\theta})}$, where T is the log partition function,

$$T(\boldsymbol{\theta}) \coloneqq \log\left(\sum_{\omega \in \Omega} e^{\boldsymbol{\phi}(\omega)\cdot\boldsymbol{\theta}}\right). \tag{4}$$

For the log partition function, $\nabla T(\boldsymbol{\theta})$ coincides with the expectation of $\boldsymbol{\phi}$ under $p(\cdot; \boldsymbol{\theta})$, and $\operatorname{dom} T^* = \operatorname{conv}\{\boldsymbol{\phi}(\omega) : \omega \in \Omega\}$. The parameter $\boldsymbol{\theta} \in R^K$ is referred to as the natural parameter of the distribution. All relevant information about an outcome ω is captured through $\boldsymbol{\phi}(\omega)$,

which is referred to as the sufficient statistic. In our setting, $\boldsymbol{\phi}(\omega)$ will be the security payoff function as defined in Section 2. Exponential family distributions are both analytically convenient and widely used. Commonly studied examples of exponential family distributions include the binomial, Poisson, Gaussian, Dirichlet, and exponential distributions Barndorff-Nielsen [1982].

As in Abernethy et al. [2014], we assume that each trader i has exponential family beliefs with natural parameter $\tilde{\boldsymbol{\theta}}_i$. From the perspective of trader i, the expected payoffs of the K market securities can then be expressed as the vector $\tilde{\boldsymbol{\mu}}_i$ with $\tilde{\mu}_{i,k} \coloneqq \sum_{\omega \in \Omega} \phi_k(\omega) p(\omega; \tilde{\boldsymbol{\theta}}_i)$.

Since the analysis of error requires some notion of ground truth, we additionally assume that the true outcome is drawn from an exponential-family distribution with natural parameter $\boldsymbol{\theta}^{\text{true}}$. The corresponding vector of expected security payoffs is denoted $\boldsymbol{\mu}^{\text{true}}$. The relationship between the traders' beliefs and ground truth comes into play only in our discussion of sampling error, where we assume that the traders' beliefs are independent and that each $\tilde{\boldsymbol{\theta}}_i$ is an unbiased sample of $\boldsymbol{\theta}^{\text{true}}$.

Following Abernethy et al. [2014], we additionally assume that traders are risk averse with exponential utility for wealth. Formally, the utility of trader i for wealth W can be written as

$$u_i(W) = -(1/a_i)e^{-a_i W}$$

where a_i is the trader's Arrow-Pratt coefficient of absolute risk aversion. Larger values of a_i imply higher levels of risk aversion. We assume that the traders' risk aversion coefficients are fixed.

Using the definitions of the expected utility U_i, the exponential family distribution $p(\cdot; \tilde{\boldsymbol{\theta}}_i)$, the log partition function T, and the exponential utility u_i, it is straightforward to show that the expected utility of trader i can be written

$$U_i(\boldsymbol{r}_i, c_i) = -\frac{1}{a_i} e^{-T(\tilde{\boldsymbol{\theta}}_i) - a_i c_i} \sum_{\omega \in \Omega} e^{\boldsymbol{\phi}(\omega) \cdot (\tilde{\boldsymbol{\theta}}_i - a_i \boldsymbol{r}_i)} = -\frac{1}{a_i} e^{T(\tilde{\boldsymbol{\theta}}_i - a_i \boldsymbol{r}_i) - T(\tilde{\boldsymbol{\theta}}_i) - a_i c_i}. \tag{5}$$

Similar expressions were used in the analysis of Abernethy et al. [2014].

4.2 Potential Functions and Uniqueness of Equilibrium Prices

Under this restricted trader model, we can construct potential functions which yield alternative characterizations of the market-clearing equilibria and the market-maker equilibria as solutions of minimization problems. These characterizations enable us to prove uniqueness of the equilibrium prices and play a crucial role in our analysis of bias and convergence error.

Consider first a market-clearing equilibrium. For each trader i, define

$$F_i(\boldsymbol{s}) \coloneqq \frac{1}{a_i} T(\tilde{\boldsymbol{\theta}}_i + a_i \boldsymbol{s}) \ .$$

From Equation (5) we can observe that $-F_i(-\boldsymbol{r}_i) + c_i$ is a monotone transformation of trader i's utility. Since each trader's utility is locally maximized at a market-clearing equilibrium, the sum of traders' utilities is also locally maximized, as is $\sum_{i=1}^{N}(-F_i(-\boldsymbol{r}_i) + c_i)$. Since the equilibrium conditions require that $\sum_{i=1}^{N} c_i = 0$, this implies that the security allocation associated with any market-clearing equilibrium is a local minimum of $\sum_{i=1}^{N} F_i(-\boldsymbol{r}_i)$.

This idea is formalized in the following theorem, the proof of which follows from an analysis of the KKT conditions of the equilibrium definition. From the theorem, it is easy to see that there are infinitely many equilibrium allocations of cash, and in general there can also be infinitely many equilibrium allocations of securities. However, by analyzing the duals of the potential, we can show

that the equilibrium prices are unique. Our proofs for this section use similar techniques to those in Frongillo and Reid [2015], and are in the appendix (with all omitted proofs).

Theorem 4.1. *Under the exponential trader model, $(\bar{\boldsymbol{r}}, \bar{\boldsymbol{c}}, \bar{\boldsymbol{\mu}})$ is a market-clearing equilibrium if and only if*

$$\bar{\boldsymbol{r}} \in \operatorname*{argmin}_{\boldsymbol{r}:\, \sum_{i=1}^N \boldsymbol{r}_i = \boldsymbol{0}} \sum_{i=1}^N F_i(-\boldsymbol{r}_i), \qquad \sum_{i=1}^N \bar{\boldsymbol{c}}_i = 0, \qquad \text{and} \qquad \bar{\boldsymbol{\mu}} = \nabla T(\tilde{\boldsymbol{\theta}}_i - a_i \bar{\boldsymbol{r}}_i) \;\; \forall i \in [N].$$

A market-clearing equilibrium always exists. Furthermore, for any market-clearing equilibrium, the equilibrium prices are unique solutions of the following dual problem:

$$\bar{\boldsymbol{\mu}} = \operatorname*{argmin}_{\boldsymbol{\mu} \in \mathbb{R}^K} \Big[\sum_i F_i^*(\boldsymbol{\mu}) \Big] \;. \tag{6}$$

Using a similar argument, we can show that the allocation associated with any market-maker equilibrium is a local minimum of the function

$$F(\boldsymbol{r}) \coloneqq \sum_{i=1}^N F_i(-\boldsymbol{r}_i) + C_b\Big(\sum_{i=1}^N \boldsymbol{r}_i \Big) \;. \tag{7}$$

Theorem 4.2. *Under the exponential trader model, $(\boldsymbol{r}^\star, \boldsymbol{c}^\star, \boldsymbol{\mu}^\star)$ is a market-maker equilibrium for cost function C_b if and only if, for the market state $\boldsymbol{s}^\star = \sum_{i=1}^N \boldsymbol{r}_i^\star$,*

$$\boldsymbol{r}^\star \in \operatorname*{argmin}_{\boldsymbol{r}} F(\boldsymbol{r}), \quad \sum_{i=1}^N \boldsymbol{c}_i^\star = C_b(\boldsymbol{0}) - C_b(\boldsymbol{s}^\star), \quad \text{and} \quad \boldsymbol{\mu}^\star = \nabla C_b(\boldsymbol{s}^\star) = \nabla T(\tilde{\boldsymbol{\theta}}_i - a_i \boldsymbol{r}_i^\star) \;\forall i \in [N].$$

A market-maker equilibrium always exists. Furthermore, for any market-maker equilibrium, the equilibrium prices are unique solutions of the following dual problem:

$$\boldsymbol{\mu}^\star = \operatorname*{argmin}_{\boldsymbol{\mu} \in \mathbb{R}^K} \Big[\sum_i F_i^*(\boldsymbol{\mu}) + bC^*(\boldsymbol{\mu}) \Big] \;. \tag{8}$$

5 Sampling Error and the Risk-Aversion Effect

The market's forecasting ability is fundamentally limited by the information present among the population of traders, and the traders' risk attitudes in communicating their information via trades. In this section, we quantify these sources of error by analyzing the discrepancy $\|\boldsymbol{\mu}^{\text{true}} - \bar{\boldsymbol{\mu}}\|$ between the true expected security values and the market-clearing equilibrium prices; here, $\|\cdot\|$ denotes the standard Euclidean norm. The following characterization of $\bar{\boldsymbol{\mu}}$ allows us to bound this discrepancy.

Theorem 5.1. *Under the exponential trader model, the unique market-clearing equilibrium prices $\bar{\boldsymbol{\mu}}$ can be written as*

$$\bar{\boldsymbol{\mu}} = \nabla T \left(\frac{\sum_{i=1}^N \tilde{\boldsymbol{\theta}}_i / a_i}{\sum_{i=1}^N 1/a_i} \right) \;. \tag{9}$$

This form reflects two possible sources of error. First, the beliefs $\tilde{\boldsymbol{\theta}}_i$ are typically noisy signals of the ground truth. Second, beliefs are weighted according to risk aversions a_i, which can skew the

prices. To formalize the latter concept, we write $N_{\text{eff}} = \left(\sum_i a_i^{-1}\right)^2 / \left(\sum_i a_i^{-2}\right)$ to denote the *effective sample size* of the weighted average. When risk aversion coefficients are equal across agents, we have $N_{\text{eff}} = N$, and when one agent has much smaller risk aversion than the others, $N_{\text{eff}} \to 1$. As the next result shows, the magnitude of the sampling error depends on the effective sample size as it relates to the number of securities and the variance in trader beliefs.

Theorem 5.2. *Under the exponential trader model, assume that the beliefs $\tilde{\boldsymbol{\theta}}_i$ are drawn independently for each trader $i \in [N]$ with mean $\mathbb{E}[\tilde{\boldsymbol{\theta}}_i] = \boldsymbol{\theta}^{\text{true}}$ and covariance $\mathbb{V}(\tilde{\boldsymbol{\theta}}_i) \preceq \sigma^2 I_K$ for some $\sigma^2 \geq 0$. For any $\delta \in (0,1)$, the market-clearing prices $\bar{\boldsymbol{\mu}}$ satisfy, with probability at least $1-\delta$, $\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}^{\text{true}}\| \leq O\left(\sigma\sqrt{K/(N_{\text{eff}}\,\delta)}\right)$. Furthermore, assuming that each a_i lies in a bounded range $(a_{\min}, a_{\max})$ where $a_{\min}, a_{\max} > 0$, we have that $N_{\text{eff}} \to \infty$ as $N \to \infty$.*

To interpret this result, note that the characterization of market-clearing prices in (6) evokes standard statistical estimators such as the maximum likelihood estimate (after rescaling the objective by $1/N$). Theorem 5.2 implies that as the number of traders grows large, the market prices $\bar{\boldsymbol{\mu}}$ converge to the ground truth $\boldsymbol{\mu}^{\text{true}}$ in probability. It is important to note that this relies on Theorem 5.1, which is an artifact of exponential utility—for general utilities, market-clearing prices may not be consistent in the statistical sense, and this extra discrepancy would need to be quantified.

For a finite number of traders, the bound in Theorem 5.2 increases with the belief variance and the number of securities, as one would expect. It decreases with the effective sample size: the information incorporated into market-clearing prices improves when risk aversions are more uniform, and when the number of agents increases.

6 Market-maker Bias

In this section, we analyze the market-maker bias, the difference between the marker-maker equilibrium prices $\boldsymbol{\mu}^\star$ and market-clearing prices $\bar{\boldsymbol{\mu}}$. We provide both global and local bounds to characterize the dependence on the liquidity b and cost function C. Our global bound shows that the bias is not more than cb for a constant c that depends on C but not on the trader beliefs. Our tighter local analysis shows that for an appropriate vector $\boldsymbol{u}$, which depends on both $\bar{\boldsymbol{\mu}}$ and C, we have $\boldsymbol{\mu}^\star = \bar{\boldsymbol{\mu}} + b\boldsymbol{u} + \boldsymbol{\varepsilon}_b$ where $\|\boldsymbol{\varepsilon}_b\| = O(b^2)$. We provide an example that shows that, under certain conditions, independent markets have up to twice as much bias as LMSR for a fixed value of b.

6.1 Smoothness and Strict Convexity Assumptions

Parts of our analysis require assumptions that go beyond convexity and differentiability of the cost function (and hence potential). We call the class of functions that satisfy them *convex$^+$ functions*. Before we define these functions, we introduce the *gradient space*, which plays a role in the structure of convex functions.[1] After stating the definitions, we provide a few examples.

Definition 6.1. Let $f : \mathbb{R}^d \to \mathbb{R}$ be a differentiable convex function. Its *gradient space*, $\mathcal{G}(f) \coloneqq \operatorname{span}\{\nabla f(\boldsymbol{u}) - \nabla f(\boldsymbol{u}') : \boldsymbol{u}, \boldsymbol{u}' \in \mathbb{R}^d\}$, is the linear space parallel to the affine hull of its set of gradients.

[1] The definitions of gradient space and convex$^+$ functions are presented here for the special case of convex functions that are finite everywhere. More general definitions that apply to functions which may equal ∞ are given in Appendix A.3.

It can be checked that if P is a projection on $\mathcal{G}(f)$ then there exists a unique point $\boldsymbol{a} \in \mathcal{G}(f)^\perp$, i.e., in the orthogonal complement of $\mathcal{G}(f)$, such that $f(\boldsymbol{u}) = f(P\boldsymbol{u}) + \boldsymbol{a}^\intercal \boldsymbol{u}$, so f is up to a linear term fully described by its values on $\mathcal{G}(f)$.

Definition 6.2. We say that a function $f : \mathbb{R}^d \to \mathbb{R}$ is *convex*$^+$ if it satisfies the following conditions: (1) f is convex; (2) f has continuous third derivatives; and (3) $\text{range}(\nabla^2 f(\boldsymbol{u})) = \mathcal{G}(f)$ for all $\boldsymbol{u} \in \mathbb{R}^d$.

The condition on the range of the Hessian ensures that f is strictly convex over $\mathcal{G}(f)$, so its gradient map is invertible over $\mathcal{G}(f)$. As a consequence of the invertibility, the Hessian of f can be expressed as a function of the gradient, i.e., there exists a matrix-valued function H_f such that $\nabla^2 f(\boldsymbol{u}) = H_f(\nabla f(u))$ (see Proposition A.8). The LMSR, the sum of independent LMSRs (IND), and the log partition function are all convex$^+$.

Example 6.3 (LMSR as a convex$^+$ function)**.** For LMSR, the gradient space is parallel to the simplex: $\mathcal{G}(C) = \{\boldsymbol{u} : \mathbf{1}^\intercal \boldsymbol{u} = 0\}$ where $\mathbf{1}$ is the all-ones-vector. The gradients of C are points in the relative interior of the simplex. Given such a point $\boldsymbol{\mu} = \nabla C(\boldsymbol{s})$, the corresponding Hessian is $\nabla^2 C(\boldsymbol{s}) = H_C(\boldsymbol{\mu}) = (\text{diag}_{k\in[K]}\, \mu_k) - \boldsymbol{\mu}\boldsymbol{\mu}^\intercal$. Here, $\text{diag}_{k\in[K]}\, \mu_k$ denotes the diagonal $K \times K$ matrix with values μ_k on the diagonal. The matrix $H_C(\boldsymbol{\mu})$ can be interpreted as the covariance matrix of the multinomial distribution $\boldsymbol{\mu}$. The null space of the matrix is exactly $\mathcal{G}(C)^\perp = \{c\mathbf{1} : c \in \mathbb{R}\}$. This means that C is constant in the all-ones direction, but strictly convex in directions from $\mathcal{G}(C)$. Finally, note that $C(\boldsymbol{s}) = C(P\boldsymbol{s}) + \boldsymbol{a}^\intercal \boldsymbol{u}$ where $P = I_K - \mathbf{1}\mathbf{1}^\intercal/K$ and $\boldsymbol{a} = \mathbf{1}/K$.

Example 6.4 (IND as a convex$^+$ function)**.** For IND, the gradient space is $\mathbb{R}^K$ and the gradients are the points in $(0, 1)^K$. In this case, $H_C(\boldsymbol{\mu}) = \text{diag}_k[\mu_k(1 - \mu_k)]$, corresponding to the covariance matrix of K independent Bernoulli variables with expectations μ_k. This matrix has full rank.

In addition to (or instead of) convexity$^+$, some of our results require Lipschitz gradients or, dually, strong convexity. To be precise, we say that a differentiable function $f : \mathbb{R}^d \to \mathbb{R}$ has a Lipschitz gradient if there exists a constant L such that $\|\nabla f(\boldsymbol{u}) - \nabla f(\boldsymbol{v})\| \le L\|\boldsymbol{u} - \boldsymbol{v}\|$ for all $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^d$, where $\|\cdot\|$ is the standard Euclidean norm. If f is twice differentiable, it suffices to check that $\nabla^2 f(\boldsymbol{u}) \preceq LI_d$ for all $\boldsymbol{u}$, where $I_d \in \mathbb{R}^{d\times d}$ is the identity. Both LMSR and IND have Lipschitz gradients since $H_C(\boldsymbol{\mu}) \preceq I_K$ in both cases. Similarly the log partition function has Lipschitz gradient, because $\nabla^2 T(\boldsymbol{\theta})$ is the covariance matrix of a bounded random variable, namely $\boldsymbol{\phi}(\omega)$ when $\omega \sim p(\omega; \boldsymbol{\theta})$.

For a convex $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$, the *subdifferential* of f at $\boldsymbol{u}$ is defined as $\partial f(\boldsymbol{u}) := \{\boldsymbol{v} \in \mathbb{R}^d : \ f(\boldsymbol{u}') \ge f(\boldsymbol{u}) + \boldsymbol{v}^\intercal(\boldsymbol{u}' - \boldsymbol{u}), \forall \boldsymbol{u}' \in \mathbb{R}^d\}$. When the subdifferential is a singleton, it coincides with the gradient. We say that f is strongly convex with the strong convexity constant σ if $f(\boldsymbol{v}) \ge f(\boldsymbol{u}) + \boldsymbol{g}^\intercal(\boldsymbol{v} - \boldsymbol{u}) + \frac{1}{2}\sigma\|\boldsymbol{v} - \boldsymbol{u}\|^2$, for all $\boldsymbol{v}, \boldsymbol{u} \in \mathbb{R}^d$ and $\boldsymbol{g} \in \partial f(\boldsymbol{u})$. A standard convex analysis result states that if $f : \mathbb{R}^d \to \mathbb{R}$ has a gradient with Lipschitz constant L then f^* is strongly convex with the strong convexity constant $\sigma = 1/L$ (see Prop. 12.60 of Rockafellar and Wets, 2009).

6.2 Bounding the Market-Maker Bias

To bound the bias, our main tool is the dual characterizations of $\boldsymbol{\mu}^\star$ and $\bar{\boldsymbol{\mu}}$ provided by Theorems 4.1 and 4.2. The two objectives in Eqs. (6) and (8) only differ in the term $bC^*(\boldsymbol{\mu})$. While the two optimization problems are formally over $\boldsymbol{\mu} \in \mathbb{R}^K$, the functions F_i^* are equal to ∞ outside the set $\mathcal{M}$ of coherent payoff prices, so the optimization is effectively only over $\mathcal{M}$. Within this set, the

function C^* is bounded and functions F_i^* are strictly convex. These two properties suffice to show that $\boldsymbol{\mu}^\star \to \bar{\boldsymbol{\mu}}$ as $b \to 0$, since the additive effect of $bC^*(\boldsymbol{\mu})$ goes to zero and strict convexity prevents the minimum from moving too far from $\bar{\boldsymbol{\mu}}$ when perturbations of the objective are small.

However, since gradients ∇F_i are Lipschitz, the functions F_i^* are not just strictly convex, but actually *strongly convex* over $\mathcal{M}$, so the sum $\sum_i F_i^*(\boldsymbol{\mu})$ is bounded below on $\mathcal{M}$ by a quadratic function of the form $F_i^*(\bar{\boldsymbol{\mu}}) + \frac{1}{2}\sigma\|\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}\|^2$ for some strong convexity constant σ, because $\bar{\boldsymbol{\mu}}$ minimizes $\sum_i F_i^*(\boldsymbol{\mu})$. At the same time, $C^*(\boldsymbol{\mu})$ is bounded below by a linear function of the form $C^*(\bar{\boldsymbol{\mu}}) + \boldsymbol{u}^\intercal(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}})$, because C^* is subdifferentiable at $\bar{\boldsymbol{\mu}}$. After some algebra, we then obtain the bound $O(b)$ on the rate of convergence $\boldsymbol{\mu}^\star \to \bar{\boldsymbol{\mu}}$ as $b \to 0$.

Theorem 6.5. *Under the exponential trader model, for any C, there exists a constant c such that $\|\boldsymbol{\mu}^\star(b; C) - \bar{\boldsymbol{\mu}}\| \le cb$ for all $b \ge 0$.*

6.3 Tighter Analysis for Convex$^+$ Cost Functions

The analysis in Theorem 6.5 is based on the strong convexity constants of F_i^* that are valid over the entire set $\mathcal{M}$. Depending on the value of $\bar{\boldsymbol{\mu}}$, these constants can be overly conservative. Moreover, the theorem only proves an upper bound, which cannot be used to compare different cost function families, since the bound may be loose for some families and tighter for others.

Our next theorem carries out a tighter local analysis, based on the properties of F_i^* and C^* at $\bar{\boldsymbol{\mu}}$, and shows that for an appropriate vector $\boldsymbol{u}$, which depends on $\bar{\boldsymbol{\mu}}$ and C, we have $\boldsymbol{\mu}^\star(b; C) = \bar{\boldsymbol{\mu}} + b\boldsymbol{u} + \boldsymbol{\varepsilon}_b$ where $\|\boldsymbol{\varepsilon}_b\| = O(b^2)$. Here, the $O(\cdot)$ is taken as $b \to 0$, so the error term $\boldsymbol{\varepsilon}_b$ goes to zero faster than the term $b\boldsymbol{u}$, which we call the *asymptotic bias vector*, and its magnitude the *asymptotic bias*. Our analysis is *local* in the sense that the constants hiding within $O(\cdot)$ may depend on $\bar{\boldsymbol{\mu}}$. In contrast with Theorem 6.5, the local analysis fully uncovers the main asymptotic term and therefore allows comparison of cost families. In our experiments in Section 8, we will see that the asymptotic bias is a highly accurate estimate of the actual bias for even moderately large values of b. This tighter analysis requires that the cost function be convex$^+$.

Theorem 6.6. *Assume that the cost function C is convex$^+$. Then the difference between the market-maker equilibrium price $\boldsymbol{\mu}^\star(b; C)$ and the market-clearing equilibrium price $\bar{\boldsymbol{\mu}}$ is*

$$\boldsymbol{\mu}^\star(b; C) = \bar{\boldsymbol{\mu}} + b\left(-\frac{\bar{a}}{N} H_T(\bar{\boldsymbol{\mu}})\partial C^*(\bar{\boldsymbol{\mu}})\right) + \boldsymbol{\varepsilon}_b \ , \quad \text{where } \|\boldsymbol{\varepsilon}_b\| = O(b^2).$$

In the expression above, $\bar{a} = N/(\sum_{i=1}^N 1/a_i)$ is the harmonic mean of risk-aversion coefficients and $H_T(\bar{\boldsymbol{\mu}})\partial C^*(\bar{\boldsymbol{\mu}})$ is guaranteed to consist of a single point even when $\partial C^*(\bar{\boldsymbol{\mu}})$ is a set.

This theorem gives us a formula to calculate the asymptotic bias for any cost function if we know the value of $\bar{\boldsymbol{\mu}}$. It can also be used to evaluate the worst-case bias against some set of possible market-clearing prices. In the following example, we work out a comparison between `LMSR` and `IND`, which does not require such knowledge. We study the case when the asymptotic bias vector under `IND` is a multiple of the asymptotic bias vector under `LMSR`, which occurs when the components of $\bar{\boldsymbol{\mu}}$ take on only two distinct values.

Example 6.7 (Bias of `LMSR` vs `IND`)**.** Assume that the components of $\bar{\boldsymbol{\mu}}$ take on only two distinct values $\mu^{(1)} > \mu^{(2)}$. We begin by evaluating $\partial C^*(\bar{\boldsymbol{\mu}})$. Let $\boldsymbol{s}^{\texttt{LMSR}}$ and $\boldsymbol{s}^{\texttt{IND}}$ denote any elements of

$\partial C^*(\bar{\boldsymbol{\mu}})$ for the two costs. From the properties of conjugates, $\boldsymbol{s} \in \partial C^*(\bar{\boldsymbol{\mu}})$ iff $\nabla C(\boldsymbol{s}) = \bar{\boldsymbol{\mu}}$, so we can obtain $\boldsymbol{s}^{\texttt{LMSR}}$ and $\boldsymbol{s}^{\texttt{IND}}$ by inverting the gradients of LMSR and IND:

$$s_k^{\texttt{LMSR}} = \log \bar{\mu}_k \ , \quad s_k^{\texttt{IND}} = \log\left(\tfrac{\bar{\mu}_k}{1-\bar{\mu}_k}\right) \quad \text{for all } k \in [K].$$

To evaluate $H_T(\bar{\boldsymbol{\mu}})\boldsymbol{s}$, note that the range of $H_T(\bar{\boldsymbol{\mu}})$ excludes the all-ones direction, so we can add any $c\mathbf{1}$ to $\boldsymbol{s}$ without affecting the value of $H_T(\bar{\boldsymbol{\mu}})\boldsymbol{s}$. We shift $\boldsymbol{s}^{\texttt{LMSR}}$ and $\boldsymbol{s}^{\texttt{IND}}$ to make the components corresponding to $\mu^{(2)}$ equal to zero. The components corresponding to $\mu^{(1)}$ then become

$$\hat{s}_k^{\texttt{LMSR}} = \log\left(\tfrac{\mu^{(1)}}{\mu^{(2)}}\right) \ , \quad \hat{s}_k^{\texttt{IND}} = \log\left(\tfrac{\mu^{(1)}}{\mu^{(2)}} \cdot \tfrac{1-\mu^{(2)}}{1-\mu^{(1)}}\right) \quad \text{for all } k \text{ such that } \bar{\mu}_k = \mu^{(1)}.$$

Thus, $\hat{\boldsymbol{s}}^{\texttt{IND}} = \eta \hat{\boldsymbol{s}}^{\texttt{LMSR}}$ where

$$\eta \coloneqq \log\left(\tfrac{\mu^{(1)}}{\mu^{(2)}} \cdot \tfrac{1-\mu^{(2)}}{1-\mu^{(1)}}\right) \Big/ \log\left(\tfrac{\mu^{(1)}}{\mu^{(2)}}\right) = 1 + \log\left(\tfrac{1-\mu^{(2)}}{1-\mu^{(1)}}\right) \Big/ \log\left(\tfrac{\mu^{(1)}}{\mu^{(2)}}\right) = 1 + \log\left(\tfrac{c+\mu^{(1)}}{c+\mu^{(2)}}\right) \Big/ \log\left(\tfrac{\mu^{(1)}}{\mu^{(2)}}\right) \ ,$$

where $c \coloneqq 1 - \mu^{(1)} - \mu^{(2)} \geq 0$. Setting $\boldsymbol{u} = (-\bar{a}/N) H_T(\bar{\boldsymbol{\mu}}) \hat{\boldsymbol{s}}^{\texttt{LMSR}}$, and ignoring the terms on the order $O(b^2)$, we therefore obtain

$$\boldsymbol{\mu}^\star(b; \texttt{LMSR}) \approx \bar{\boldsymbol{\mu}} + b\boldsymbol{u} \ , \qquad \boldsymbol{\mu}^\star(b; \texttt{IND}) \approx \bar{\boldsymbol{\mu}} + \eta b \boldsymbol{u} \ .$$

Thus, for the same value of b, the bias of IND is larger by a factor of η than the bias of LMSR. Since $\mu^{(1)} > \mu^{(2)}$ and $c \geq 0$, we see that $1 < \eta \leq 2$. The value $\eta = 2$ is attained when $c = 0$, which occurs whenever $K = 2$, because then $\mu^{(1)} + \mu^{(2)} = 1$.

7 Convergence Error

We now turn to the convergence error, the difference between the prices $\boldsymbol{\mu}^t$ at round t and the marker-maker equilibrium prices $\boldsymbol{\mu}^\star$. We consider two types of trader dynamics, both of which allow the market to be interpreted as a convex optimization algorithm performing block-coordinate descent. Frongillo and Reid [2015] showed a global convergence bound of $O(1/t)$ after t rounds of trade. We prove a tighter local bound of $O(\gamma^t)$ with $\gamma = 1 - \Omega(b)$. For our *all-securities* dynamics, we give a lower bound on γ, fully characterizing the convergence rate. Along the way, we prove a local convergence bound for block-coordinate descent, extending results of Nesterov [2012].

7.1 Trader Dynamics

To study convergence properties of the market, we need to posit a model of how the traders arrive in the market and which securities they buy or sell, as a function of their current holdings of securities $\boldsymbol{r}_i$ and cash c_i, and the current market state $\boldsymbol{s}$. We refer to such a model as *trader dynamics*. We consider two simple trader dynamics:

- *All-securities dynamics* (ASD). In each round, a trader $i \in [N]$ is chosen uniformly at random. This trader then buys a bundle $\boldsymbol{\delta} \in \mathbb{R}^K$ which optimizes her utility, i.e., if the current state of the market is $\boldsymbol{s}$ and the current cash and security allocations of the trader are c_i and $\boldsymbol{r}_i$, then the trader picks $\boldsymbol{\delta}$ maximizing $U_i(\boldsymbol{r}_i + \boldsymbol{\delta},\, c_i - C_b(\boldsymbol{s} + \boldsymbol{\delta}) + C_b(\boldsymbol{s}))$.

- *Single-security dynamics* (SSD). In each round, a trader $i \in [N]$ is chosen uniformly at random and this trader picks a security $k \in [K]$ uniformly at random. The trader then buys a quantity $\delta \in \mathbb{R}$ of the kth security to optimize her utility. Let $\boldsymbol{e}_k$ be the kth vector of the standard basis. Then the trader picks δ maximizing $U_i(\boldsymbol{r}_i + \delta \boldsymbol{e}_k,\, c_i - C_b(\boldsymbol{s} + \delta \boldsymbol{e}_k) + C_b(\boldsymbol{s}))$.

The all-securities dynamics has been recently studied by Frongillo and Reid [2015]. This model assumes that traders are able to calculate the optimal bundle over all securities, which may not be a realistic assumption when the number of securities is large. The single-security dynamics, in which we only assume that traders can optimize over a single security at a time, is more appropriate for computationally limited traders.

We formalize both dynamics in a unified analysis via *blocks*. Specifically, we assume that the coordinates $[N] \times [K]$ are partitioned into blocks $\alpha \in \mathcal{A}$, where $\alpha \subseteq [N] \times [K]$, such that $\biguplus_{\alpha \in \mathcal{A}} \alpha = [N] \times [K]$. Blocks are disjoint subsets of coordinates of the overall allocation vector $\boldsymbol{r} \in \mathbb{R}^{NK}$. We further assume that each block $\alpha \in \mathcal{A}$ is fully contained within coordinates corresponding to some trader i. Thus each block α can be written as $\{i\} \times \beta$ for some $i \in [N]$ and $\beta \subseteq [K]$. For ASD, we have N blocks $\mathcal{A} = \big\{\{i\} \times [K] : i \in [N]\big\}$. For SSD, we have NK blocks $\mathcal{A} = \big\{\{i\} \times \{k\} : i \in [N], k \in [K]\big\}$.

Let $E_\alpha \in \mathbb{R}^{NK \times |\alpha|}$ be the *embedding matrix* for the block α. It maps $|\alpha|$-dimensional vectors $\boldsymbol{u} \in \mathbb{R}^{|\alpha|}$ to vectors $\boldsymbol{v} \in \mathbb{R}^{NK}$ that are zero everywhere except for the block α, where they coincide with $\boldsymbol{u}$. The range of E_α is exactly the set of vectors that are zero outside the block α. Its transpose $E_\alpha^\intercal$ projects vectors from $\mathbb{R}^{NK}$ to $\mathbb{R}^{|\alpha|}$ by removing all coordinates outside the block α. For any subset $\beta \subseteq [K]$, we similarly define the embedding matrices $E_\beta \in \mathbb{R}^{K \times |\beta|}$.

The next theorem shows that optimizing utility, as in ASD or SSD, corresponds to the greedy optimization of the potential function F along the coordinates of the corresponding block:

Theorem 7.1. *Let $\alpha = \{i\} \times \beta$, where $\beta \subseteq [K]$, be a block of coordinates controlled by trader i. Assume that trader i has security bundle $\boldsymbol{r}_i$ and c_i units of cash, and the current market state is $\boldsymbol{s}$. Then*

$$\operatorname*{argmax}_{\boldsymbol{\delta} \in \operatorname{range}(E_\beta)} U_i(\boldsymbol{r}_i + \boldsymbol{\delta},\, c_i - C_b(\boldsymbol{s} + \boldsymbol{\delta}) + C_b(\boldsymbol{s})) = \operatorname*{argmin}_{\boldsymbol{\delta} \in \operatorname{range}(E_\beta)} F(\boldsymbol{r}_{-i},\, \boldsymbol{r}_i + \boldsymbol{\delta}) \ ,$$

where $\boldsymbol{r}_{-i}$ denotes the concatenation of $\boldsymbol{r}_j$ across $j \neq i$.

Since both dynamics correspond to variants of randomized block-coordinate descent on the potential F, we can appeal to tools from convex optimization to analyze convergence error.

7.2 Global Convergence Rate

Let $\boldsymbol{r}^\star$ denote an arbitrary minimizer of F and let $F^\star$ denote the minimum value of F. Also, let $\boldsymbol{r}^t$ denote the allocation vector and $\boldsymbol{\mu}^t$ the market price after the tth trade. We next show that if C has a Lipschitz gradient then for bounding the convergence error $\|\boldsymbol{\mu}^t - \boldsymbol{\mu}^\star\|$ it suffices to bound the suboptimality of the current objective value, $F(\boldsymbol{r}^t) - F^\star$:

Theorem 7.2. *If ∇C has the Lipschitz constant L_C then $F(\boldsymbol{r}^t) - F^\star \geq b \cdot \frac{1}{2L_C} \|\boldsymbol{\mu}^t - \boldsymbol{\mu}^\star\|^2$.*

As noted by Frongillo and Reid [2015], for C with a Lipschitz gradient we can use the results of Nesterov [2012] to obtain an $O(1/t)$ bound on $\mathbb{E}\big[F(\boldsymbol{r}^t)\big] - F^\star$, where the expectation is over the choices of traders and securities at each round. The Markov inequality then yields the bound $F(\boldsymbol{r}^t) - F^\star \leq O(1/\rho t)$ with probability at least $1 - \rho$. In the remainder of this section we show that we can significantly improve on this bound and obtain a local linear convergence, meaning

$\mathbb{E}\left[F(\boldsymbol{r}^t)\right] - F^\star \leq O(\gamma^t)$ for some $\gamma < 1$, once $\boldsymbol{r}^t$ is sufficiently close to any minimizer $\boldsymbol{r}^\star$. In our simulations we show that the local-regime bound is accurate already for moderate values of t.

Both Nesterov [2012] and Frongillo and Reid [2015] show linear convergence when the objective F is globally strongly convex, which is not the case for our potential under LMSR or IND cost. On the other hand, our analysis only requires local strong convexity structure and we obtain both upper and lower bounds on the rate of convergence γ. We begin by modifying the existing block-coordinate descent analysis of Nesterov [2012], and then focus on the specific optimization problem arising in our setting.

7.3 Local Convergence Rate of Block-coordinate Descent

In this section, we consider general unconstrained convex minimization, but use the same notation as in the rest of the paper. The key difference from the standard analysis of Nesterov [2012] is the focus on the local convergence in the neighborhood of the solution, rather than global convergence. This analysis is not specific to our setting, and may be of independent interest.

We consider the optimization problem $\min_{\boldsymbol{r}\in\mathbb{R}^{NK}} F(\boldsymbol{r})$, where $F : \mathbb{R}^{NK} \to \mathbb{R}$ is a differentiable convex function bounded below. We are given a set of blocks $\alpha \in \mathcal{A}$, which partition the coordinates $[N]\times[K]$. The *block-coordinate descent* algorithm sets the initial iterate $\boldsymbol{r}^0 = \mathbf{0}$, and in each iteration chooses an index $\alpha \in \mathcal{A}$ uniformly at random and fully optimizes the objective over the coordinates in block α: given the current iterate $\boldsymbol{r}^t$, the new iterate is $\boldsymbol{r}^{t+1} = \operatorname{argmin}_{\boldsymbol{r}'\in\boldsymbol{r}^t+\operatorname{range}(E_\alpha)} F(\boldsymbol{r}')$.

While the objective in our setting is not globally strongly convex, it is strongly convex locally, so the optimization eventually stays within a region where the strong convexity constant is bounded away from zero, yielding a linear convergence rate.

Definition 7.3. We say that γ_{high} is an *upper bound on the local convergence rate* of an algorithm if, with probability 1 under the randomness of the algorithm, the algorithm reaches an iteration t_0 such that for some $c > 0$ and all $t \geq t_0$, $\mathbb{E}\left[F(\boldsymbol{r}^t) \,\middle|\, \boldsymbol{r}^{t_0}\right] - F^\star \leq c\gamma_{\text{high}}^{t-t_0}$.

In order to compare families of cost functions, it is also useful to consider lower bounds:

Definition 7.4. We say that γ_{low} is a *lower bound on the local convergence rate* if $\gamma_{\text{high}} \geq \gamma_{\text{low}}$ holds for all upper bounds γ_{high}.

We use properties of F in the neighborhood of $\boldsymbol{r}^\star$ to derive bounds on the local convergence rate. For the lower bound, the theorem assumes that the optimization problem is *non-degenerate* with respect to the algorithm in the sense that the probability of reaching an iterate $\boldsymbol{r}^t$ which attains a minimum is zero; in other words, only approximate solutions are reached in a finite number of steps. This condition holds for our potential function under `ASD`, whenever there are at least three agents with distinct beliefs.

To state the theorem we need to introduce some additional notation. For a matrix $M \in \mathbb{R}^{NK\times NK}$, let $M_{\alpha,\alpha}$ denote the block consisting of rows and columns in α, i.e., $M_{\alpha,\alpha} \coloneqq E_\alpha^\intercal M E_\alpha$. And let $\mathcal{D} : \mathbb{R}^{NK\times NK} \to \mathbb{R}^{NK\times NK}$ be the operation of retaining only the block diagonal of a matrix, i.e., $\mathcal{D}(M) \coloneqq \operatorname{diag}_{\alpha\in\mathcal{A}} M_{\alpha,\alpha}$. Let M^+ denote the pseudoinverse of M. Finally, let $\boldsymbol{g}^t \coloneqq \nabla F(\boldsymbol{r}^t)$ denote the gradient of the objective in the tth iteration. Similarly to the local analysis of market-maker bias, the local analysis of convergence rate requires that the objective be convex$^+$:

Theorem 7.5. *Let F be a convex$^+$ function attaining a minimum at $\boldsymbol{r}^\star$. Let $H^\star := \nabla^2 F(\boldsymbol{r}^\star)$ and assume there exists $\lambda > F^\star$ and non-negative constants σ_{low}, σ_{high}, and ℓ such that*

$$\sigma_{\text{low}} \leq \frac{\mathbb{E}\left[\ (\boldsymbol{g}^t)^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{g}^t \ \middle|\ \boldsymbol{r}^{t-\ell}\ \right]}{\mathbb{E}\left[\ (\boldsymbol{g}^t)^\intercal (H^\star)^+ \boldsymbol{g}^t \ \middle|\ \boldsymbol{r}^{t-\ell}\ \right]} \leq \sigma_{\text{high}} \tag{10}$$

whenever $F(\boldsymbol{r}^{t-\ell}) \leq \lambda$, where the expectation is over the choice of updates by the algorithm. Then, for all $\varepsilon > 0$, the local convergence rate is bounded above by $\gamma_{\text{high}} = 1 - \sigma_{\text{low}}/|\mathcal{A}| + \varepsilon$. If the optimization problem is non-degenerate, the local convergence is also bounded below by $\gamma_{\text{low}} = 1 - \sigma_{\text{high}}/|\mathcal{A}|$.

The ratio bounded in Eq. (10) can be interpreted as a curvature of the quadratic form $\mathcal{D}(H^\star)^+$ under the norm described by the quadratic form $(H^\star)^+$. Larger values of the ratio (larger curvature) mean faster convergence. We will refer to the bounds σ_{low} and σ_{high} as *lower and upper bounds on local strong convexity* of F (under randomized block-coordinate descent updates). Any non-trivial lower bound, i.e., $\sigma_{\text{low}} > 0$, yields local linear convergence rate since it implies $\gamma_{\text{high}} < 1$, since ε can be chosen arbitrarily small.

If we know the Hessian $H^\star$, we can obtain a simple lower bound σ_{low} and a linear convergence by considering all directions in the span of gradients (and setting $\ell = 0$). The span of gradients coincides with $\mathcal{G}(F)$, because $\mathbf{0}$ is a valid gradient (since F attains a minimum). Thus, we can obtain σ_{low} by the following generalized eigenvalue calculation:

$$\sigma_{\text{low}} = \min_{\boldsymbol{u} \in \mathcal{G}(F) \setminus \{\mathbf{0}\}} \frac{\boldsymbol{u}^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{u}}{\boldsymbol{u}^\intercal (H^\star)^+ \boldsymbol{u}} = \lambda_{\min}\big((H^\star)^{1/2} \mathcal{D}(H^\star)^+ (H^\star)^{1/2}\big)\ , \tag{11}$$

where $\lambda_{\min}(\cdot)$ is the smallest positive eigenvalue. This is a valid setting of σ_{low}, because for all $\boldsymbol{u} \in \mathcal{G}(F)$ we then have $\sigma_{\text{low}} \boldsymbol{u}^\intercal (H^\star)^+ \boldsymbol{u} \leq \boldsymbol{u}^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{u}$ and therefore, when $F(\boldsymbol{r}^{t-\ell}) \leq \lambda$, also

$$\sigma_{\text{low}} \mathbb{E}\left[\ (\boldsymbol{g}^t)^\intercal (H^\star)^+ \boldsymbol{g}^t \ \middle|\ \boldsymbol{r}^{t-\ell}\ \right] \leq \mathbb{E}\left[\ (\boldsymbol{g}^t)^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{g}^t \ \middle|\ \boldsymbol{r}^{t-\ell}\ \right].$$

This value of σ_{low} is non-zero, because $\mathcal{G}(F) = \operatorname{range}(H^\star)^+ \subseteq \operatorname{range} \mathcal{D}(H^\star)^+$. We pursue this style of analysis for single-securities dynamics, where we derive σ_{low} using Eq. (11), but do so directly in terms of Hessians of functions C and T rather than the potential F. For all-securities dynamics, we consider $\ell > 0$ and take advantage of the averaging effect of expectations in Eq. (10), which yields a tighter lower bound σ_{low} and a non-trivial upper bound σ_{high}.

7.4 Local Convergence Rate of the Market

We now apply Theorem 7.5 to our setting to show linear local convergence for the potential function F in terms of Hessians of cost C and log partition T. We highlight key results here. Appendix E.6 contains additional details.

As in the local analysis of bias, we assume that C is convex$^+$. We also use the notation $D := \operatorname{diag}_{i \in [N]} a_i$ and $P := I_N - \mathbf{1}\mathbf{1}^\intercal / N$. The matrix P is the projection matrix on the set of centered vectors, i.e., vectors $\boldsymbol{u}$ in $\mathbb{R}^N$ such that $\mathbf{1}^\intercal \boldsymbol{u} = 0$. We use $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ to denote the smallest and the largest positive eigenvalue of a matrix.

Theorem 7.6. *Assume that C is convex$^+$. Let $H_T := H_T(\bar{\boldsymbol{\mu}})$, $H_C := H_C(\bar{\boldsymbol{\mu}})$, and D_C be the diagonal matrix with the diagonal of H_C. For all-securities dynamics, local strong convexity is bounded from below and above by*

$$\sigma_{\text{low}}^{\texttt{ASD}} = 2b \cdot \lambda_{\min}(PDP) \cdot \lambda_{\min}\big(H_T^{1/2} H_C^{+} H_T^{1/2} \big) - O(b^2) \ ,$$
$$\sigma_{\text{high}}^{\texttt{ASD}} = 2b \cdot \lambda_{\max}(PDP) \cdot \lambda_{\max}\big(H_T^{1/2} H_C^{+} H_T^{1/2} \big) + O(b^2) \ .$$

For single-security dynamics, local strong convexity is bounded from below by

$$\sigma_{\text{low}}^{\texttt{SSD}} = b \cdot \lambda_{\min}(PDP) \cdot \lambda_{\min}\big(H_T^{1/2} D_C^{+} H_T^{1/2} \big) - O(b^2) \ .$$

For `ASD`, Theorem 7.6 proves linear convergence with the rate $\gamma = 1 - \Theta(b)$. This means that the convergence gets worse as $b \to 0$, leading to a trade-off with the bias, which decreases as $b \to 0$. As we will see in numerical experiments, these bounds on the convergence rate are quite tight.

For `SSD`, we only present a lower bound, which suffices to establish a linear convergence rate. This bound is worse by a factor of two than the bound for `ASD`. This we believe is only an artifact of a looser analysis and we conjecture that the bounds we proved for `ASD` also hold for `SSD`, with H_C replaced by D_C. Our numerical experiments in Appendix F.2 suggest this might indeed be the case.

Example 7.7 (Convergence of `LMSR` under `ASD`)**.** We next demonstrate the tightness of our bounds for `ASD`. Consider the setting when $N \geq 2$ and the risk aversion of all traders equals a. Then $PDP = aPI_NP = aP$, and since P is a non-zero projection matrix, we obtain $\lambda_{\min}(PDP) = a\lambda_{\min}(P) = a$ and similarly $\lambda_{\max}(PDP) = a$. Furthermore, if the cost is `LMSR` then $H_C = H_T$, and thus $H_T^{1/2} H_C^{+} H_T^{1/2} = I_K$. Therefore, Theorem 7.6 yields the bounds $\sigma_{\text{low}}^{\texttt{ASD}} = 2ab - O(b^2)$ and $\sigma_{\text{high}}^{\texttt{ASD}} = 2ab + O(b^2)$, whose main asymptotic terms match exactly. Thus, the objective decreases at the rate γ^t and the convergence error at the rate $\gamma^{t/2}$ with $\gamma = 1 - 2ab/N + O(b^2)$, and the linear term in b cannot be improved.

8 Numerical experiments

In this section, we present numerical simulations to evaluate the tightness of our theoretical bounds. We consider a setting in which there is a complete market over $K = 5$ securities. We simulate $N = 10$ traders with exponential utilities and exponential-family beliefs. These values of N and K are large enough to demonstrate the tightness of our results, while small enough that simulations are tractable. We assume that all traders $i \in [N]$ have the same risk aversion coefficients $a_i = 1$. We fix the ground truth likelihoods of each outcome (or equivalently, ground truth expected security payoffs) $\boldsymbol{\mu}^{\text{true}}$ and corresponding natural parameters $\boldsymbol{\theta}^{\text{true}}$, and generate the natural parameters $\tilde{\boldsymbol{\theta}}_i$ of each trader i's beliefs as a (single) sample from Normal$(\boldsymbol{\theta}^{\text{true}}, \sigma^2 I_K)$. We consider two variants:

- *Uniform Beliefs*: All outcomes are equally likely. We set $\boldsymbol{\theta}^{\text{true}} = \mathbf{0}$ and $\sigma = 1$.
- *Single-Peaked Beliefs*: One outcome is more likely than the others. Here we set $\theta_1^{\text{true}} = \log(1 - \nu \cdot (k-1))$ and $\theta_k^{\text{true}} = \log(\nu)$ for $k \neq 1$. We use $\nu = 0.02$ and $\sigma = 5$.

Figure 1 shows the trader beliefs and market-clearing equilibrium prices from Eq. (9).

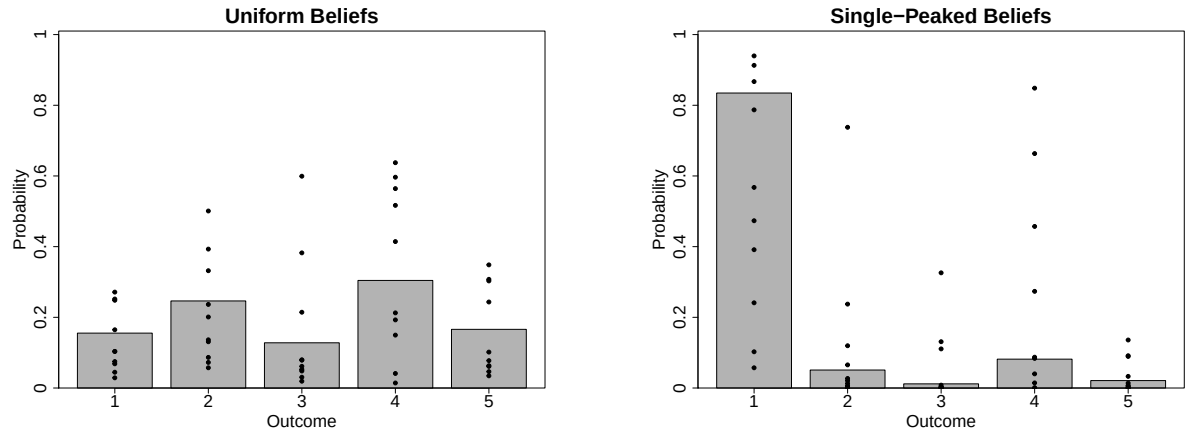

Figure 1: Two sets of beliefs of the $N = 10$ traders about the $K = 5$ outcomes. The beliefs were sampled once and then fixed in all experiments. The gray bars show the market-clearing equilibrium prices $\bar{\boldsymbol{\mu}}$ as in Eq. (9).

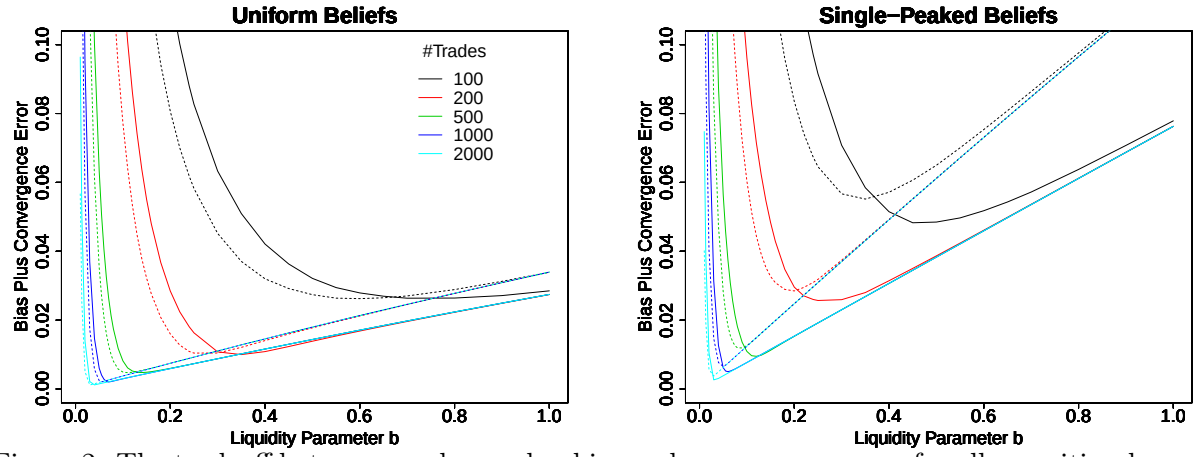

Figure 2: The tradeoff between marker-maker bias and convergence error for all-securities dynamics with different beliefs and cost functions. Solid lines show the total bias and convergence error of the LMSR after various numbers of trades, averaged over 20 random trade sequences. Dotted lines show the same for IND.

Bias/Convergence Tradeoffs. We first examine the tradeoff that arises between market-maker bias and convergence error as the liquidity parameter of the market is adjusted. Since our main interest is in the effect of the cost function C and liquidity parameter b on error, we ignore the sources of error that do not depend on the choice of cost function. Figure 2 shows the combined bias and convergence error, $\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}^t(b; C)\|_1$, as a function of liquidity, for different beliefs and cost functions under all-security dynamics after different numbers of trades have occurred. (Other choices of norm lead to similar results.) The results for single-security dynamics are similar, and are reported in Appendix F.2. The minimum point on each curve tells us the optimal value of the liquidity parameter b for the particular setting and number of trades. When the market has not been running long, larger values of b lead to lower error. On the other hand, smaller values of b are preferable as the number of trades grows, with the combined error approaching 0 for small b. The combined error of the LMSR is similar to that of the sum of independent LMSRs (IND) under uniform beliefs, but the LMSR produces lower combined error for single-peaked beliefs.

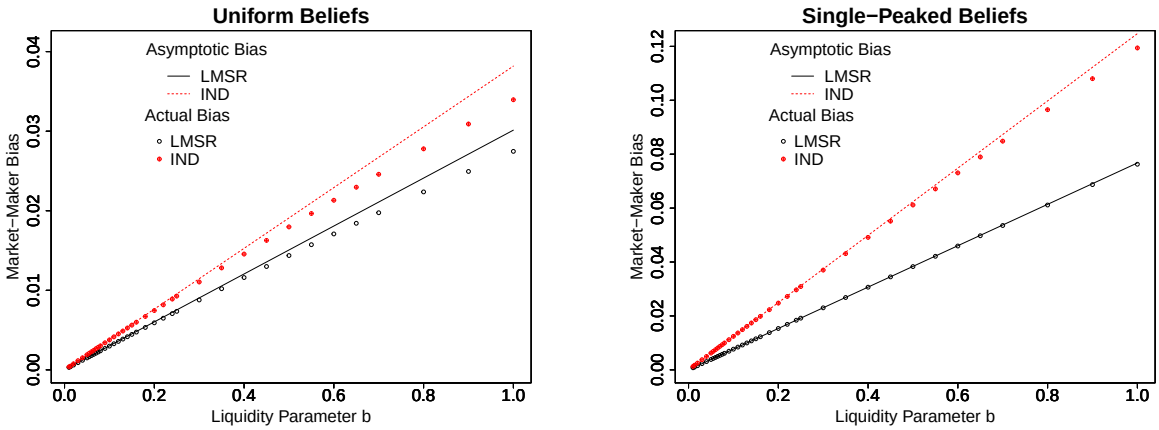

Figure 3: Market-maker bias as a function of b for different beliefs and cost functions.

Market-Maker Bias We next zoom in on the market-maker bias to empirically evaluate our bounds from Section 6. From Theorem 6.6, we know that $\|\boldsymbol{\mu}^\star(b;C) - \bar{\boldsymbol{\mu}}\|_1 \approx b \cdot \frac{\bar{a}}{N} \|H_T(\bar{\boldsymbol{\mu}})\partial C^*(\bar{\boldsymbol{\mu}})\|_1$. In Figure 3, we plot the empirical bias $\|\boldsymbol{\mu}^\star(b,C) - \bar{\boldsymbol{\mu}}\|_1$ as a function of b for both LMSR and IND under uniform and single-peaked beliefs, and in each case compare this bias with the approximation implied by the theory. We find that although Theorem 6.6 only gives an asymptotic guarantee as $b \to 0$, the approximation above is fairly accurate even for moderate values of b. As we showed for a stylized example in Section 6, for any fixed value of b, the bias of IND is higher than the bias of LMSR. We see a greater difference for single-peaked beliefs.

Convergence Error Finally, we turn to the convergence error. We first show that the local linear convergence rate kicks in very quickly—essentially from the start of trade in our simulations. We then examine the tightness of our bounds on the local convergence rate from Section 7.4.

From Theorem 7.2, we know that $F(\boldsymbol{r}^t) - F^\star$ is an upper bound on $\|\boldsymbol{\mu}^t - \boldsymbol{\mu}^\star\|^2$, where we suppress the implicit dependence on C and b. Rather than examining the convergence of prices directly, we examine convergence of the objective, which will be more convenient in the discussion below. Figure 4 shows the empirical value of $\hat{\mathbb{E}}\left[F(\boldsymbol{r}^t)\right] - F^\star$, where the expectation is the empirical average over the 20 random sequences, as a function of the number of trades, plotted on a log scale, for our two belief sets and cost functions under the all-securities trade dynamics. The results for single-security dynamics are similar, and are included in Appendix F.2. In all settings, the log of convergence error appears linear, matching the local asymptotic analysis in Section 7.4. In other words, there exist some $\hat{c}$ and $\hat{\gamma}$ such that, empirically, we have for all t, $\hat{\mathbb{E}}\left[F(\boldsymbol{r}^t)\right] - F^\star \approx \hat{c}\hat{\gamma}^t$.

To examine the tightness of the bounds from Section 7.4, we dig more deeply into the value of this empirical constant $\hat{\gamma}$, which depends on C and b. Since this approximation holds for any sufficiently large t, we can define $\hat{\gamma}$ by choosing some t_1 and t_2 and setting $\hat{\gamma} = \left(\frac{\hat{\mathbb{E}}\left[F(\boldsymbol{r}^{t_2})\right] - F^\star}{\hat{\mathbb{E}}\left[F(\boldsymbol{r}^{t_1})\right] - F^\star}\right)^{1/(t_2 - t_1)}$. If $\hat{\gamma}$ is the correct asymptotic convergence rate, then from Theorem 7.5, we should have $1 - \sigma_{\text{high}}/N \leq \hat{\gamma} \leq 1 - \sigma_{\text{low}}/N$ for values of σ_{high} and σ_{low} that satisfy Eq. (10), since $|\mathcal{A}| = N$ for the all-securities dynamics. Rearranging terms, we would expect that, for sufficiently large t_1 and t_2,

$$\sigma_{\text{low}} \leq N\left(1 - \left(\frac{\hat{\mathbb{E}}\left[F(\boldsymbol{r}^{t_2})\right] - F^\star}{\hat{\mathbb{E}}\left[F(\boldsymbol{r}^{t_1})\right] - F^\star}\right)^{1/(t_2 - t_1)}\right) \leq \sigma_{\text{high}}. \tag{12}$$

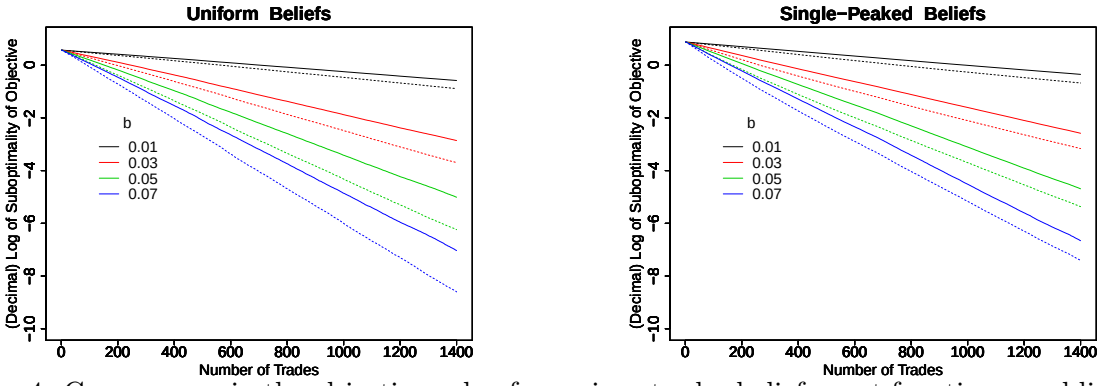

Figure 4: Convergence in the objective value for various trader beliefs, cost functions, and liquidity parameters under the all-securities dynamics. Solid lines show the log error in objective for LMSR, dotted lines for IND.

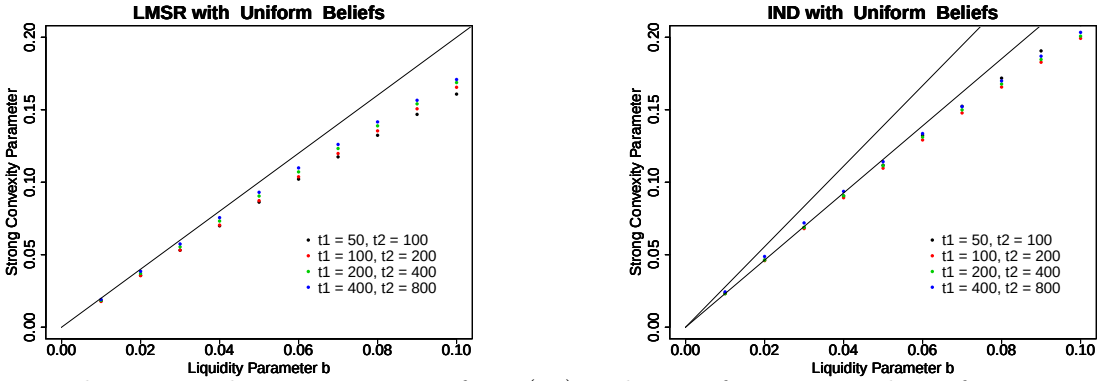
LMSR with Uniform Beliefs
IND with Uniform Beliefs
Strong Convexity Parameter
Liquidity Parameter b
t1 = 50, t2 = 100
t1 = 100, t2 = 200
t1 = 200, t2 = 400
t1 = 400, t2 = 800

Figure 5: The empirical strong convexity from (12) under ASD for various values of t_1, t_2 represented as dots, asymptotic bounds for σ_{low} and σ_{high} from Theorem 7.6 (ignoring $O(b^2)$ terms) represented as lines.

We refer to this quantity that is upper and lower bounded by σ_{high} and σ_{low} as the *empirical strong convexity* $\hat{\sigma}$. Note that σ_{low} implicitly depends on b and C.

We can now check how well our theoretical lower and upper bounds on local strong convexity bound the empirical strong convexity $\hat{\sigma}$. In Figure 5, we plot $\hat{\sigma}$ as a function of b using different values of t_1 and t_2 and compare it with the asymptotic bounds of σ_{low} and σ_{high} computed as in Theorem 7.6, dropping the terms that are $O(b^2)$. We would expect to see $\sigma_{\text{low}} \leq \hat{\sigma} \leq \sigma_{\text{high}}$ as b goes to 0, and indeed this is the case. For the LMSR, the values of σ_{high} and σ_{low} coincide, and the empirical values for $\hat{\sigma}$ agree for small b. For space, we only present the plots for uniform beliefs in Figure 5, but we present the plots for single-peaked beliefs in Appendix F.1.

We have illustrated that our local characterizations of both bias and convergence error are tight, which justifies the use of the local analysis. Our theoretical framework therefore provides a meaningful way to compare the detailed quantitative error tradeoffs inherent in different choices of cost functions and liquidity levels.

References

Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. 2013. Efficient Market Making via Convex Optimization, and a Connection to Online Learning. *ACM Transactions on Economics and Computation* 1, 2 (2013), Article 12.

Jacob Abernethy, Sindhu Kutty, Sébastien Lahaie, and Rahul Sami. 2014. Information Aggregation in Exponential Family Markets. In *15th ACM Conference on Economics and Computation (EC)*.

Ole Barndorff-Nielsen. 1982. *Exponential Families*. Wiley Online Library.

Joyce Berg, Robert Forsythe, Forrest Nelson, and Thomas Rietz. 2008. Results from a dozen years of election futures markets research. *Handbook of experimental economics results* 1 (2008), 742–751.

Yiling Chen and David M. Pennock. 2007. A Utility Framework for Bounded-Loss Market Makers. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI)*.

Miroslav Dudík, Sébastien Lahaie, David M. Pennock, and David Rothschild. 2013. A combinatorial prediction market for the US elections. In *Proceedings of the fourteenth ACM conference on Electronic Commerce*. 341–358.

Rafael Frongillo and Mark D. Reid. 2015. Convergence Analysis of Prediction Markets via Randomized Subspace Descent. In *Advances in Neural Information Processing Systems (NIPS)*.

Robin Hanson. 2003. Combinatorial Information Market Design. *Information Systems Frontiers* 5, 1 (2003), 105–119.

Yu. Nesterov. 2012. Efficiency of Coordinate Descent Methods on Huge-Scale Optimization Problems. *SIAM Journal on Optimization* 22, 2 (2012), 341–362.

K. B. Petersen and M. S. Pedersen. 2012. The Matrix Cookbook. (Nov 2012). Version 20121115.

R. Tyrrell Rockafellar. 1970. *Convex analysis*. Princeton University Press.

R. Tyrrell Rockafellar and Roger J-B Wets. 2009. *Variational analysis*. Springer-Verlag.

David Rothschild. 2009. Forecasting elections: comparing prediction markets, polls, and their biases. *Public Opinion Quarterly* 73, 5 (2009), 895–916.

Christian Slamka, Bernd Skiera, and Martin Spann. 2013. Prediction market performance and market liquidity: A comparison of automated market makers. *IEEE Transactions on Engineering Management* 60, 1 (2013), 169–185.

Justin Wolfers and Eric Zitzewitz. 2004. Prediction markets. *The Journal of Economic Perspectives* 18, 2 (2004), 107–126.

A Mathematical background

A.1 Vectors, matrices, intervals

Vectors are denoted by small boldface italics $\boldsymbol{u}, \boldsymbol{\mu}, \ldots$, matrices by large italics $A, B, P, \ldots$. The norm $\|\cdot\|$ denotes the standard Euclidean norm for vectors, and the operator norm for matrices, i.e., $\|A\| = \sup_{\boldsymbol{u}:\, \|\boldsymbol{u}\|=1} \|A\boldsymbol{u}\|$. The *range of a matrix*, denoted $\mathrm{range}(\cdot)$, is the span of its columns. For a symmetric matrix A, its pseudoinverse, denoted A^+, is the unique symmetric matrix with $\mathrm{range}(A^+) = \mathrm{range}(A)$ such that $A^+A = AA^+ = P$ where P is the projection on $\mathrm{range}(A)$.

For symmetric matrices A and B, we write $A \preceq B$ to denote that $B - A$ is positive-semidefinite. We use the notation $A \simeq B \pm C$ to denote $B - C \preceq A \preceq B + C$. For scalars, we similarly write $a \simeq b \pm c$ to mean $a \in [b - c, b + c]$.

We use $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ to denote the smallest and the largest positive eigenvalue of a symmetric positive-semidefinite matrix. We also write $\lambda_{\min}(A, B)$ and $\lambda_{\max}(A, B)$ for the generalized minimum and maximum eigenvalues defined as follows whenever $B \neq \mathbf{0}$:

$$\lambda_{\min}(A, B) \coloneqq \min_{\boldsymbol{u} \in \mathrm{range}(B) \setminus \{\mathbf{0}\}} \frac{\boldsymbol{u}^\intercal A \boldsymbol{u}}{\boldsymbol{u}^\intercal B \boldsymbol{u}} \ , \qquad \lambda_{\max}(A, B) \coloneqq \max_{\boldsymbol{u} \in \mathrm{range}(B) \setminus \{\mathbf{0}\}} \frac{\boldsymbol{u}^\intercal A \boldsymbol{u}}{\boldsymbol{u}^\intercal B \boldsymbol{u}} \ .$$

The following formulas follow by substituting $\boldsymbol{v} = B^{1/2}\boldsymbol{u}$:

$$\lambda_{\min}(A, B) = \lambda_{\min}\big((B^{1/2})^+ A (B^{1/2})^+ \big) \quad \text{if } \mathrm{range}(A) \subseteq \mathrm{range}(B),$$
$$\lambda_{\max}(A, B) = \lambda_{\max}\big((B^{1/2})^+ A (B^{1/2})^+ \big) \quad \text{if } \mathrm{range}(A) \supseteq \mathrm{range}(B).$$

By Eq. (180) of Petersen and Pedersen [2012], we also have for any matrix A (not necessarily square):

$$\lambda_{\min}(AA^\intercal) = \lambda_{\min}(A^\intercal A) \ , \quad \lambda_{\max}(AA^\intercal) = \lambda_{\max}(A^\intercal A) \ . \tag{13}$$

The following two results relate A and its pseudoinverse A^+:

Proposition A.1. *Let A and B be symmetric matrices such that* $\mathrm{range}(A) = \mathrm{range}(B)$. *Then $A \preceq B$ if and only if $B^+ \preceq A^+$.*

Proof. Assume that $\mathrm{range}(A) = \mathrm{range}(B) = \mathcal{L}$. Then we have the following equivalences

$$\begin{aligned}
A \preceq B \ \text{ iff } \ & \boldsymbol{u}^\intercal A \boldsymbol{u} \leq \boldsymbol{u}^\intercal B \boldsymbol{u} \text{ for all } \boldsymbol{u} \in \mathcal{L} \setminus \{\mathbf{0}\} \\
\text{iff } \ & 1 \leq \min_{\boldsymbol{u} \in \mathcal{L} \setminus \{\mathbf{0}\}} \frac{\boldsymbol{u}^\intercal B \boldsymbol{u}}{\boldsymbol{u}^\intercal A \boldsymbol{u}} \\
\text{iff } \ & 1 \leq \lambda_{\min}(B, A) = \lambda_{\min}\big((A^{1/2})^+ B (A^{1/2})^+ \big) \\
& \qquad = \lambda_{\min}\big(B^{1/2} A^+ B^{1/2} \big) = \lambda_{\min}(A^+, B^+) \qquad \text{(by Eq. 13)} \\
\text{iff } \ & 1 \leq \min_{\boldsymbol{u} \in \mathcal{L} \setminus \{\mathbf{0}\}} \frac{\boldsymbol{u}^\intercal A^+ \boldsymbol{u}}{\boldsymbol{u}^\intercal B^+ \boldsymbol{u}} \\
\text{iff } \ & B^+ \preceq A^+
\end{aligned}$$

□

Proposition A.2 (Blockwise Inversion)**.** *Let $A, B, C \in \mathbb{R}^{d \times d}$, where A and C are symmetric, and let $S = C - B^\intercal A^+ B$. Assume the following conditions hold:*

- $\mathrm{range}(A) \cap \mathrm{range}(C) = \{\mathbf{0}\}$.

- $\mathrm{range}(B) \subseteq \mathrm{range}(A)$.
- $\mathrm{range}(B^\intercal) \subseteq \mathrm{range}(C)$.
- $\mathrm{range}(S) = \mathrm{range}(C)$.

Then

$$(A + B + B^\intercal + C)^+ = A^+ + (I - A^+B)S^+(I - A^+B)^\intercal \ .$$

Proof. This follows by the formula in Section 9.1.3 of Petersen and Pedersen [2012] for the block matrix

$$\begin{pmatrix} U^\intercal AU & U^\intercal BV \\ V^\intercal BU & V^\intercal CV \end{pmatrix}$$

where U and V are matrices with an orthonormal basis of $\mathrm{range}(A)$ and $\mathrm{range}(C)$, respectively. □

The final result of this section relates the optimization of a quadratic form with a matrix A to the quadratic form with the matrix A^+:

Proposition A.3. *Let $A \in \mathbb{R}^{d\times d}$ be a symmetric positive-semidefinite matrix and let $\boldsymbol{u} \in \mathrm{range}(A)$. Then*

$$\min_{\boldsymbol{\delta}\in\mathbb{R}^d} \left(\boldsymbol{\delta}^\intercal \boldsymbol{u} + \frac{1}{2}\boldsymbol{\delta}^\intercal A \boldsymbol{\delta} \right) = -\frac{1}{2}\boldsymbol{u}^\intercal A^+ \boldsymbol{u} \ .$$

Proof. The result follows from convex conjugacy of quadratic convex functions, see page 108 of Rockafellar [1970]. □

A.2 Convex analysis

If $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ then the *epigraph* of f is the set of points in $\mathbb{R}^d \times \mathbb{R}$ that lie on or above the graph of f. The function f is called *convex* if its epigraph is convex. It is called *closed* if its epigraph is closed, and *proper* if it is not identically equal to ∞. The *effective domain* of f is the set of points where it is finite, denoted $\operatorname{dom} f$. The convex hull of a set S, written $\operatorname{conv} S$, is the smallest convex set containing S. For instance, the simplex in $\mathbb{R}^d$ is the convex hull of the vectors of standard basis.

An *affine subspace* of $\mathbb{R}^d$ is any set that can be written as $\mathcal{A} = \{\boldsymbol{a} + \boldsymbol{u} : \ \boldsymbol{u} \in \mathcal{L}\}$ for some fixed vector $\boldsymbol{a}$ where $\mathcal{L}$ is a linear subspace of $\mathbb{R}^d$. We refer to $\mathcal{L}$ as the linear space *parallel to* $\mathcal{A}$. The *affine hull* of a non-empty set S, denoted $\operatorname{aff} S$, is the smallest affine set that contains S. The *relative interior* of a set S, denoted $\operatorname{ri} S$, is the interior of S under the topology of its affine hull. The *relative boundary* of a set S consists of points in the closure of S that are not in its relative interior, $(\operatorname{cl} S)\backslash(\operatorname{ri} S)$. For instance, the affine hull of the simplex consists of vectors $\boldsymbol{u}$ such that $\boldsymbol{u}^\intercal \mathbf{1} = 1$, where $\mathbf{1}$ is the all-ones vector. The parallel linear space is $\{\boldsymbol{u} : \ \boldsymbol{u}^\intercal \mathbf{1} = 0\}$. The simplex has an empty interior, but its relative interior consists of points $\{\boldsymbol{u} \in (0,1)^d : \ \boldsymbol{u}^\intercal \mathbf{1} = 1\}$. The relative boundary of the simplex consists of those points in the simplex which have at least one coordinate equal to zero.

For a convex $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$, the *subdifferential* of f at $\boldsymbol{u}$ is defined as $\partial f(\boldsymbol{u}) := \{\boldsymbol{v} \in \mathbb{R}^d : \ f(\boldsymbol{u}') \geq f(\boldsymbol{u}) + \boldsymbol{v}^\intercal(\boldsymbol{u}' - \boldsymbol{u}), \forall \boldsymbol{u}' \in \mathbb{R}^d\}$. Any convex $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ is *subdifferentiable*, i.e., its subdifferential is non-empty, at all points in $\operatorname{ri} \operatorname{dom} f$. If the subdifferential is a singleton, it coincides with the gradient. Given a proper convex function $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$, we define its *convex conjugate* $f^* : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ by $f^*(\boldsymbol{\mu}) := \sup_{\boldsymbol{u}\in\mathbb{R}^d}[\boldsymbol{u}^\intercal \boldsymbol{\mu} - f(\boldsymbol{u})]$. For a closed proper

convex function f, its conjugate is also closed, proper and convex, and the following statements are equivalent:

$$\boldsymbol{\mu} \in \partial f(\boldsymbol{u}) \quad \text{iff} \quad \boldsymbol{u} \in \partial f^*(\boldsymbol{u}) \quad \text{iff} \quad f(\boldsymbol{u}) + f^*(\boldsymbol{\mu}) = \boldsymbol{u}^\intercal \boldsymbol{\mu} \ .$$

We will use the following variant of a duality result known as Fenchel's duality:

Theorem A.4 (Fenchel's duality)**.** *Let $f : \mathbb{R}^d \to \mathbb{R}$ and $g : \mathbb{R}^K \to \mathbb{R} \cup \{\infty\}$ be closed proper convex functions and $A \in \mathbb{R}^{K \times d}$. Assume that there exists some $\boldsymbol{u} \in \mathbb{R}^d$ such that $A\boldsymbol{u} \in \operatorname{ri}(\operatorname{dom} g)$ and some $\boldsymbol{\mu} \in \operatorname{ri}(\operatorname{dom} g^*)$ such that $A^\intercal \boldsymbol{\mu} \in \operatorname{ri}(\operatorname{dom} f^*)$. Then*

$$\inf_{\boldsymbol{u} \in \mathbb{R}^d} [f(-\boldsymbol{u}) + g(A\boldsymbol{u})] = \sup_{\boldsymbol{\mu} \in \mathbb{R}^K} [-f^*(A^\intercal \boldsymbol{\mu}) - g^*(\boldsymbol{\mu})]$$

and both the supremum and the infimum are attained. Vectors $\hat{\boldsymbol{u}}$ and $\hat{\boldsymbol{\mu}}$ are their respective solutions if and only if $A^\intercal \hat{\boldsymbol{\mu}} \in \partial f(-\hat{\boldsymbol{u}})$ and $\hat{\boldsymbol{\mu}} \in \partial g(A\hat{\boldsymbol{u}})$.

Proof. The results follows from Corollary 31.2.1 and Theorem 31.3 of Rockafellar [1970]. □

A.3 Convex$^+$ functions

Throughout the paper we work with functions that satisfy additional assumptions beyond convexity. We refer to them as *convex$^+$ functions*. They have a close relationship to functions of *Legendre type* (see Sec. 26 of Rockafellar, 1970). Before we define convex$^+$ functions, we introduce the *gradient space*, which plays a role in their structure.

Definition A.5. Let $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ be differentiable on the interior of its domain $D \coloneqq \operatorname{int} \operatorname{dom} f$. Its *gradient space*, denoted $\mathcal{G}(f)$, is the linear space parallel to the affine hull of the set of its gradients,

$$\mathcal{G}(f) \coloneqq \operatorname{span}\{\nabla f(\boldsymbol{u}) - \nabla f(\boldsymbol{u}') : \ \boldsymbol{u}, \boldsymbol{u}' \in D\} \ .$$

Definition A.6. Let $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ and $D \coloneqq \operatorname{int} \operatorname{dom} f$. We say that f is *convex$^+$* if it satisfies the following conditions:

1. f is closed and convex.
2. D is non-empty.
3. f has continuous third derivatives on D.
4. $\operatorname{range}(\nabla^2 f(\boldsymbol{u})) = \mathcal{G}(f)$ for all $\boldsymbol{u} \in D$.
5. $\lim_{t \to \infty} \|\nabla f(\boldsymbol{u}_t)\| = \infty$ whenever $\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots$ is a sequence in D converging to a boundary point of D.

Proposition A.7. *Let $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ be convex$^+$, $D \coloneqq \operatorname{int} \operatorname{dom} f$. Let P be the projection on $\mathcal{G}(f)$ and $\boldsymbol{a}$ the unique point in $\mathcal{A} \cap \mathcal{G}(f)^\perp$, where $\mathcal{A}$ is the affine hull of the set of gradients of f. Then the following statements hold:*

1. *$f(\boldsymbol{u}) = f(P\boldsymbol{u}) + \boldsymbol{a}^\intercal \boldsymbol{u}$.*
2. *$\nabla f(\boldsymbol{u}) = \nabla f(P\boldsymbol{u})$.*
3. *$\nabla^2 f(\boldsymbol{u}) = \nabla^2 f(P\boldsymbol{u})$.*
4. *f is strictly convex on $D \cap \mathcal{G}(f)$.*
5. *∇f is one-to-one on $D \cap \mathcal{G}(f)$.*

Proof. We prove the first statement by the Mean Value Theorem. First, since f is differentiable on an open convex D, it must be actually continuously differentiable on D (by Corollary 25.5.1 of Rockafellar, 1970), so the Mean Value Theorem can be applied on D. Let $\boldsymbol{u} \in D$ and $\boldsymbol{v} \in \mathcal{G}(f)^\perp$. Then for any $\boldsymbol{u}' = \boldsymbol{u} + t\boldsymbol{v} \in D$, we have, for some $\bar{\boldsymbol{u}}$ on the line segment connecting $\boldsymbol{u}$ and $\boldsymbol{u}'$,

$$f(\boldsymbol{u}') = f(\boldsymbol{u}) + [\nabla f(\bar{\boldsymbol{u}})]^\intercal t\boldsymbol{v} = f(\boldsymbol{u}) + t\boldsymbol{a}^\intercal \boldsymbol{v} \tag{14}$$

where the second equality follows because $\boldsymbol{v} \perp \mathcal{G}(f)$. We argue that the entire line $\{\boldsymbol{u} + t\boldsymbol{v} : t \in \mathbb{R}\}$ must be contained in D. For contradiction, assume it intersects the boundary of D at $\boldsymbol{u}^\star = \boldsymbol{u} + t^\star \boldsymbol{v}$, and say $t^\star > 0$. Consider an increasing sequence $0 = t_1, t_2, \ldots$ converging to $t^\star$. Eq. (14) holds for $\boldsymbol{u}'$ replaced by $\boldsymbol{u}_i = \boldsymbol{u} + t_i \boldsymbol{v}$ as well as points $\tilde{\boldsymbol{u}}_i = \tilde{\boldsymbol{u}} + t_i \boldsymbol{v}$, where $\tilde{\boldsymbol{u}}$ is in a small enough neighborhood of $\boldsymbol{u}$ along directions in $\mathcal{G}(f)$. This means that $\nabla f(\boldsymbol{u}_i) = P\nabla f(\boldsymbol{u}) + \boldsymbol{a} = \nabla f(\boldsymbol{u})$. However, this is not possible for convex$^+$ functions, because the norms of their gradients go to ∞ towards the boundary. Similar argument holds for $t^\star < 0$. This means that the entire line $\{\boldsymbol{u} + t\boldsymbol{v} : t \in \mathbb{R}\}$ must be in D. This holds for arbitrary $\boldsymbol{u} \in D$, so D can be written as $D = D_0 + \mathcal{G}(f)^\perp$ where $D_0 \subseteq \mathcal{G}(f)$. Eq. (14) now implies that statement (1) holds over D. Since f is closed, the statement also holds over $\operatorname{dom} f$, which then necessarily has form $\operatorname{dom} f = S_0 + \mathcal{G}(f)^\perp$ where $S_0 \subseteq \mathcal{G}(f)$. Therefore, statement (1) also holds for $\boldsymbol{u} \notin \operatorname{dom} f$.

The remaining statements are more straightforward. To prove the second statement, note that for any $\boldsymbol{u}' \in D$, its gradient can be decomposed as $\nabla f(\boldsymbol{u}') = P\nabla f(\boldsymbol{u}') + \boldsymbol{a}$. Thus, $\nabla f(\boldsymbol{u}) = P\nabla f(P\boldsymbol{u}) + \boldsymbol{a} = \nabla f(P\boldsymbol{u})$. For the third statement, we have $\nabla^2 f(\boldsymbol{u}) = P[\nabla^2 f(P\boldsymbol{u})]P = \nabla^2 f(P\boldsymbol{u})$, because $\operatorname{range}(\nabla^2 f(P\boldsymbol{u})) = \mathcal{G}(f)$. The fourth statement is equivalent to the fifth statement and they follow because $\operatorname{range}(\nabla^2 f(P\boldsymbol{u})) = \mathcal{G}(f)$. □

This proposition immediately implies that the Hessian of f can be expressed as a function of the gradient of f. We will denote such a function H_f:

Proposition A.8. *Let $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ be convex$^+$ and $D \coloneqq \operatorname{int} \operatorname{dom} f$. Let $D' \coloneqq \{\nabla f(\boldsymbol{u}) : \boldsymbol{u} \in D\}$ be the set of its gradients. Then there exists a map $H_f : D' \to \mathbb{R}^{d \times d}$ such that $\nabla^2 f(\boldsymbol{u}) = H_f(\nabla f(\boldsymbol{u}))$ for all $\boldsymbol{u} \in D$.*

Proof. Let $D_0 = D \cap \mathcal{G}(f)$. Proposition A.7 implies that ∇f is a bijection from D_0 to D'. Denoting its inverse from D' to D_0 as $\boldsymbol{h}$, we can then define the map H_f via $H_f(\boldsymbol{\mu}) = \nabla^2 f(\boldsymbol{h}(\boldsymbol{\mu}))$. Now, for any $\boldsymbol{u} \in D$, we have

$$\nabla^2 f(\boldsymbol{u}) = \nabla^2 f(P\boldsymbol{u}) = \nabla^2 f(\boldsymbol{h}(\nabla f(P\boldsymbol{u}))) = H_f(\nabla f(P\boldsymbol{u})) = H_f(\nabla f(\boldsymbol{u})) \ . \qquad \square$$

Proposition A.9. *Let $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ be a convex$^+$ function and let $\mathcal{A} = \operatorname{aff} \operatorname{dom} f^*$. Then there exists a convex$^+$ function $g : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ such that the following hold:*

1. *g agrees with f^* on $\mathcal{A}$.*
2. *$\mathcal{G}(g)$ is parallel to $\mathcal{A}$, or equivalently $\mathcal{G}(g) = \mathcal{G}(f)$.*
3. *For $\boldsymbol{\mu} \in \operatorname{ri} \operatorname{dom} f^*$: $\partial f^*(\boldsymbol{\mu}) = \nabla g(\boldsymbol{\mu}) + \mathcal{G}(f)^\perp$.*
4. *For $\boldsymbol{\mu} \in \operatorname{ri} \operatorname{dom} f^*$: $\nabla^2 g(\boldsymbol{\mu}) = H_f^+(\boldsymbol{\mu})$.*

Proof. We begin by representing f via a function of Legendre type and then rely on the properties of such functions to obtain g. In particular, we note that by Proposition A.7, the function f is defined by its values on $\mathcal{G}(f)$, except for a linear term, and so we construct a function f_0 that is a transformation of f on $\mathcal{G}(f)$ and is of Legendre type.

To start, let $D = \operatorname{int}\operatorname{dom} f$, and $D_0 = D \cap \mathcal{G}(f)$. Assume that $\mathcal{G}(f)$ is d_0 dimensional $d_0 \leq d$, and let $A \in \mathbb{R}^{d \times d_0}$ be a matrix whose columns form an orthonormal basis of $\mathcal{G}(f)$. Then $A^\intercal A = I_{d_0}$, i.e., $A^\intercal$ is the left inverse of A. The matrix A is an injective linear map from $\mathbb{R}^{d_0} \to \mathbb{R}^d$, but it is also a bijection from $\mathbb{R}^{d_0}$ to $\mathcal{G}(f)$. We also have $AA^\intercal = P$, where P is the projection on $\mathcal{G}(f)$. Define the function $f_0 : \mathbb{R}^{d_0} \to \mathbb{R} \cup \{\infty\}$ as

$$f_0(\boldsymbol{v}) \coloneqq f(A\boldsymbol{v}) \ .$$

Let $S = \operatorname{int}\operatorname{dom} f_0$, so $S = A^\intercal D_0$. We have

$$\nabla f_0(\boldsymbol{v}) = A^\intercal \nabla f(A\boldsymbol{v}), \quad \nabla^2 f_0(\boldsymbol{v}) = A^\intercal \nabla^2 f(A\boldsymbol{v}) A, \quad \nabla^3 f_0(\boldsymbol{v})[\cdot,\cdot,\cdot] = \nabla^3 f(A\boldsymbol{v})[A(\cdot), A(\cdot), A(\cdot)],$$

so in particular f_0 has continuous third derivatives over S. We next argue that f_0 is of *Legendre type* in the sense of Rockafellar [1970], page 258. For that we need to check that it satisfies the following conditions:

(a) *S is non-empty.*
This follows, because D_0 is non-empty and $D_0 \subseteq \mathcal{G}(f)$. Now, $A^\intercal$, as a linear map, is a bijection from $\mathcal{G}(f)$ to $\mathbb{R}^{d_0}$, so the set $S = A^\intercal D_0$ is also non-empty.

(b) *f_0 is differentiable throughout S.*
Similarly to the previous property, this holds, because f is differentiable throughout D_0.

(c) *$\lim_{t\to\infty} \|\nabla f_0(\boldsymbol{v}_t)\| = \infty$ whenever $\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots$ is a sequence in S converging to a boundary point of S.*
If $\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots$ is any sequence in D_0 converging to the relative boundary of D_0, then this point is on the boundary of D and therefore $\|\nabla f(\boldsymbol{u}_t)\| \to \infty$, because f is convex$^+$. Now, suppose we are given a sequence $\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots$ in S converging to a boundary point of S. Then $\boldsymbol{u}_t = A\boldsymbol{v}_t$ is exactly a sequence in D_0 converging to the relative boundary of D_0, so $\|\nabla f(A\boldsymbol{v}_t)\| \to \infty$. Since $\nabla f(A\boldsymbol{v}_t) \in \mathcal{G}(f)$ and the row space of $A^\intercal$ coincides with $\mathcal{G}(f)$, we also have $\|\nabla f_0(\boldsymbol{v}_t)\| = \|A^\intercal \nabla f(A\boldsymbol{v}_t)\| \to \infty$.

(d) *f_0 is strictly convex on S.*
Since f_0 has a continuous Hessian on S, it suffices to shaw that its Hessian is full rank, i.e., its range is $\mathbb{R}^{d_0}$. For any $\boldsymbol{v} \in S$, we have $\operatorname{range}(\nabla^2 f_0(\boldsymbol{v})) = \operatorname{range}(A^\intercal \nabla^2 f(A\boldsymbol{v})A)$, which must be $\mathbb{R}^{d_0}$, because $\operatorname{range}(\nabla^2 f(A\boldsymbol{v})) = \mathcal{G}(f)$ and A is a bijection from $\mathbb{R}^{d_0}$ to $\mathcal{G}(f)$.

We now express f in terms of f_0. Note that by conjugacy, the affine hull of gradients of f coincides with $\mathcal{A} = \operatorname{aff}\operatorname{dom} f^*$. By Proposition A.7, the function f is defined by its values on $\mathcal{G}(f)$, except for the linear term, described by the unique $\boldsymbol{a} \in \mathcal{A} \cap \mathcal{G}^\perp$. We defined f_0 to exactly represent f on $\mathcal{G}(f)$, so for any $\boldsymbol{u}' \in \mathcal{G}(f)$, we have $f_0(A^\intercal \boldsymbol{u}') = f(AA^\intercal \boldsymbol{u}') = f(\boldsymbol{u}')$, because $AA^\intercal = P$. Thus, for any $\boldsymbol{u} \in \mathbb{R}^d$ we have, by Proposition A.7,

$$f(\boldsymbol{u}) = f(P\boldsymbol{u}) + \boldsymbol{a}^\intercal \boldsymbol{u} = f_0(A^\intercal P\boldsymbol{u}) + \boldsymbol{a}^\intercal \boldsymbol{u} = f_0(A^\intercal \boldsymbol{u}) + \boldsymbol{a}^\intercal \boldsymbol{u} \ , \tag{15}$$

because the row space of $A^\intercal$ coincides with $\mathcal{G}(f)$. This relationship between f and f_0 implies the following relationship for their conjugates (by Theorems 12.3 and 16.3 of Rockafellar, 1970):

$$f^*(\boldsymbol{\mu}) = \inf_{\boldsymbol{y} \in \mathbb{R}^d :\, A\boldsymbol{y} = \boldsymbol{\mu} - \boldsymbol{a}} f_0^*(\boldsymbol{y}) \ , \tag{16}$$

where the infimum of an empty set is ∞. The linear map A is injective and $\text{range}(A) = \mathcal{G}(f)$, so the linear system $A\boldsymbol{v} = \boldsymbol{\mu} - \boldsymbol{a}$ has a single solution $\boldsymbol{v} = A^\intercal(\boldsymbol{\mu} - \boldsymbol{a})$ when $(\boldsymbol{\mu} - \boldsymbol{a}) \in \mathcal{G}(f)$, and no solutions when $(\boldsymbol{\mu} - \boldsymbol{a}) \notin \mathcal{G}(f)$. Therefore,

$$f^*(\boldsymbol{\mu}) = \begin{cases} f_0^*\big(A^\intercal(\boldsymbol{\mu} - \boldsymbol{a})\big) & \text{if } \boldsymbol{\mu} \in \mathcal{A}, \\ \infty & \text{if } \boldsymbol{\mu} \notin \mathcal{A}. \end{cases}$$

Since f_0 is of Legendre type, so is its conjugate f_0^* (see Theorem 26.5 of Rockafellar, 1970), which means that f_0^* with $S^* \coloneqq \operatorname{int}\operatorname{dom} f_0^*$ satisfies the properties (a)–(d). The function g is constructed as follows:

$$g(\boldsymbol{\mu}) \coloneqq f_0^*\big(A^\intercal(\boldsymbol{\mu} - \boldsymbol{a})\big) = f_0^*(A^\intercal\boldsymbol{\mu}) \ ,$$

where the equality follows because $\boldsymbol{a} \in \mathcal{G}(f)^\perp$. This differs from the expression for f^* in that it does not equal to ∞ outside $\mathcal{A}$. Before we argue that g is convex$^+$, we analyze the gradient, Hessian and third derivatives of f_0^*. From the properties of the conjugates, we know that ∇f_0^* is the inverse of ∇f_0, so for all $\boldsymbol{y} \in S^*$,

$$\nabla f_0^*(\boldsymbol{y}) = [\nabla f_0]^{-1}(\boldsymbol{y}) \ .$$

Since ∇f_0 is continuously differentiable and its derivative (i.e., the Hessian of f_0) is an invertible matrix, the inverse of ∇f_0 is also continuously differentiable and its derivative (i.e., the Hessian of f_0^*) is

$$\nabla^2 f_0^*(\boldsymbol{y}) = \big[\nabla^2 f_0\big([\nabla f_0]^{-1}(\boldsymbol{y})\big)\big]^{-1} = \big[\nabla^2 f_0\big(\nabla f_0^*(\boldsymbol{y})\big)\big]^{-1} \ . \tag{17}$$

In particular note that $\nabla^2 f_0^*(\boldsymbol{y})$ is also an invertible matrix. We next show that f_0^* has a continuous third derivative, by using the chain rule to argue that $\nabla^2 f_0^*$ is continuously differentiable. This follows, because $\nabla^2 f_0^*$ is the composition of (i) the matrix inversion $[\cdot]^{-1}$, (ii) the Hessian map $\nabla^2 f_0$, and (iii) the conjugate gradient ∇f_0^*, and all of them are continuously differentiable at points where the respective derivatives are taken; specifically, the matrix inversion is taken at an invertible matrix $\nabla^2 f_0\big(\nabla f_0^*(\boldsymbol{y})\big)$, the Hessian at $\nabla f_0^*(\boldsymbol{y}) \in S$, and the conjugate gradient at $\boldsymbol{y} \in S^*$.

Now we can show that g is convex$^+$. From the definition of g, we have

$$\nabla g(\boldsymbol{\mu}) = A\nabla f_0^*(A^\intercal\boldsymbol{\mu}), \quad \nabla^2 g(\boldsymbol{\mu}) = A\nabla f_0^*(A^\intercal\boldsymbol{\mu})A^\intercal, \tag{18}$$

and

$$\nabla^3 g(\boldsymbol{\mu})[\cdot,\cdot,\cdot] = \nabla^3 f_0^*(A^\intercal\boldsymbol{\mu})[A^\intercal(\cdot), A^\intercal(\cdot), A^\intercal(\cdot)]. \tag{19}$$

Let $D^* \coloneqq \operatorname{int}\operatorname{dom} g$. From the definition of g, we have $D^* = AS^* + \mathcal{G}(f)^\perp$. Note that $\mathcal{G}(f_0^*) = \mathbb{R}^{n_0}$, because $\operatorname{dom} f_0$ has a non-empty interior. Therefore, by Eq. (18), $\mathcal{G}(g) = \text{range}(A) = \mathcal{G}(f)$. We next verify that g satisfies properties (1)–(5) of convexity$^+$, relying on the properties (a)–(d) satisfied by f_0^* and S^*:

1. *g is closed and convex.*
   This follows, because f_0^* is closed and convex.
2. *D^* is non-empty.*
   This follows, since f_0^* satisfies (a), so S^* is non-empty, and so is D^*.
3. *g has continuous third derivatives on D^*.*
   This follows by Eq. (19), because f_0^* has continuous third derivatives on S^*.
4. $\text{range}(\nabla^2 g(\boldsymbol{\mu})) = \mathcal{G}(g)$ *for all* $\boldsymbol{\mu} \in D^*$.
   Since the Hessians of f_0^* are full rank, Eq. (18) implies that $\text{range}(\nabla^2 g(\boldsymbol{\mu})) = \text{range}(A) = \mathcal{G}(g)$.

5. $\lim_{t\to\infty}\|\nabla g(\boldsymbol{\mu}_t)\| = \infty$ *whenever* $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \ldots$ *is a sequence in* D^* *converging to a boundary point of* D^*.
If $\boldsymbol{\mu}_t$ converges to a point on the border of $D^* = AS^* + \mathcal{G}(f)^\perp$, then the sequence of points $\boldsymbol{y}_t = A^\intercal \boldsymbol{\mu}_t$ converges to the border of S^*. Since f_0^* satisfies property (c), this means that $\|\nabla f_0^*(\boldsymbol{y}_t)\| \to \infty$. And since A is injective, also $\|\nabla g(\boldsymbol{\mu}_t)\| = \|A\nabla f_0^*(\boldsymbol{y}_t)\| \to \infty$.

We now prove that g has the properties stated in the theorem:

1. *g agrees with f^* on $\mathcal{A}$.*
Immediate from the definition of g.
2. *$\mathcal{G}(g)$ is parallel to $\mathcal{A}$, or equivalently $\mathcal{G}(g) = \mathcal{G}(f)$.*
As we already argued, Eq. (18) and the fact that $\mathcal{G}(f_0^*) = \mathbb{R}^{n_0}$ imply that $\mathcal{G}(g) = \text{range}(A) = \mathcal{G}(f)$.
3. *For $\boldsymbol{\mu} \in \operatorname{ri}\operatorname{dom} f^*$: $\partial f^*(\boldsymbol{\mu}) = \nabla g(\boldsymbol{\mu}) + \mathcal{G}(f)^\perp$.*
Note that $\boldsymbol{\mu} \in \mathcal{A}$, so $(\boldsymbol{\mu} - \boldsymbol{a}) \in \mathcal{G}(f)$. The statement follows by the following chain of equivalences

$$\begin{aligned}
\boldsymbol{u} \in \partial f^*(\boldsymbol{\mu}) \quad &\text{iff} \quad \nabla f(\boldsymbol{u}) = \boldsymbol{\mu} \\
&\text{iff} \quad A\nabla f_0(A^\intercal \boldsymbol{u}) + \boldsymbol{a} = \boldsymbol{\mu} && \text{(by Eq. 15)} \\
&\text{iff} \quad \nabla f_0(A^\intercal \boldsymbol{u}) = A^\intercal(\boldsymbol{\mu} - \boldsymbol{a}) = A^\intercal \boldsymbol{\mu} && \text{(because } (\boldsymbol{\mu} - \boldsymbol{a}) \in \mathcal{G}(f) = \text{range}(A)\text{)} \\
&\text{iff} \quad \nabla f_0^*(A^\intercal \boldsymbol{\mu}) = A^\intercal \boldsymbol{u} \\
&\text{iff} \quad A\nabla f_0^*(A^\intercal \boldsymbol{\mu}) = AA^\intercal \boldsymbol{u} = P\boldsymbol{u} && \text{(because } A \text{ is injective)} \\
&\text{iff} \quad \nabla g(\boldsymbol{\mu}) = P\boldsymbol{u} && \text{(by Eq. 18)} \\
&\text{iff} \quad \boldsymbol{u} \in \nabla g(\boldsymbol{\mu}) + \mathcal{G}(f)^\perp && \text{(because } \nabla g(\boldsymbol{\mu}) \in \text{range}(A) = \mathcal{G}(f)\text{)}
\end{aligned}$$

4. *For $\boldsymbol{\mu} \in \operatorname{ri}\operatorname{dom} f^*$: $\nabla^2 g(\boldsymbol{\mu}) = H_f^+(\boldsymbol{\mu})$.*
From Eqs. (18) and (17), we have

$$\begin{aligned}
\nabla^2 g(\boldsymbol{\mu}) &= A[\nabla^2 f_0^*(A^\intercal \boldsymbol{\mu})]A^\intercal \\
&= A\big[\nabla^2 f_0\big(\nabla f_0^*(A^\intercal \boldsymbol{\mu})\big)\big]^{-1} A^\intercal \ .
\end{aligned}$$

By Proposition A.8, $H_f(\boldsymbol{\mu}) = \nabla^2 f(\boldsymbol{u})$ for any $\boldsymbol{u}$ such that $\nabla f(\boldsymbol{u}) = \boldsymbol{\mu}$, which is equivalent to $\boldsymbol{u} \in \partial f^*(\boldsymbol{\mu})$. Above, we have shown that $\nabla g(\boldsymbol{\mu}) \in \partial f^*(\boldsymbol{\mu})$, so $H_f(\boldsymbol{\mu}) = \nabla^2 f(\nabla g(\boldsymbol{\mu}))$. We continue the derivation of $H_f(\boldsymbol{\mu})$ using Eqs. (15) and (18):

$$\begin{aligned}
H_f(\boldsymbol{\mu}) &= \nabla^2 f(\nabla g(\boldsymbol{\mu})) \\
&= A\big[\nabla^2 f_0\big(A^\intercal \nabla g(\boldsymbol{\mu})\big)\big]A^\intercal \\
&= A\big[\nabla^2 f_0\big(A^\intercal A\nabla f_0^*(A^\intercal \boldsymbol{\mu})\big)\big]A^\intercal \\
&= A\big[\nabla^2 f_0\big(\nabla f_0^*(A^\intercal \boldsymbol{\mu})\big)\big]A^\intercal \ ,
\end{aligned}$$

where the last equation follows, because $A^\intercal A = I_{d_0}$. Since $\mathcal{G}(g) = \mathcal{G}(f)$, the ranges of $\nabla^2 g(\boldsymbol{\mu})$ and $H_f(\boldsymbol{\mu})$ coincide with $\mathcal{G}(f)$. From the above derivations of $\nabla^2 g(\boldsymbol{\mu})$ and $H_f(\boldsymbol{\mu})$, we also have

$$[\nabla^2 g(\boldsymbol{\mu})]H_f(\boldsymbol{\mu}) = H_f(\boldsymbol{\mu})[\nabla^2 g(\boldsymbol{\mu})] = AA^\intercal = P \ ,$$

so indeed $\nabla^2 g(\boldsymbol{\mu}) = H_f^+(\boldsymbol{\mu})$. □

Thanks to the continuity of third derivatives of f and the continuity of second derivatives of g from Proposition A.9, we can easily prove a local Lipschitz property for H_f:

Proposition A.10. *Let $f : \mathbb{R}^d \to \mathbb{R} \cup \{\infty\}$ be convex$^+$ and D' be the set of its gradients (necessarily open within* aff D'*). Let $\boldsymbol{\mu} \in D'$ and let B be a closed ball (in* aff D'*) centered at $\boldsymbol{\mu}$ and fully contained in D'. Then there exists a constant c such that for all $\boldsymbol{\mu}' \in B$*

$$H_f(\boldsymbol{\mu}') \simeq \big(1 \pm c\|\boldsymbol{\mu}' - \boldsymbol{\mu}\|\big) H_f(\boldsymbol{\mu}) \ .$$

Proof. Let g be the function from Proposition A.9. Similarly as we argued in the proof of Proposition A.9, we can write

$$H_f(\boldsymbol{\mu}) = \nabla^2 f(\nabla g(\boldsymbol{\mu})) \ ,$$

because $\nabla g(\boldsymbol{\mu}) \in \partial f^*(\boldsymbol{\mu})$, and thus $\nabla f(\nabla g(\boldsymbol{\mu})) = \boldsymbol{\mu}$. Since, g has continuous second derivatives and f has continuous third derivatives, the Mean Value Theorem implies that ∇g is Lipschitz continuous on any compact subset of D', and $\nabla^2 f$ is Lipschitz continuous on any compact subset of $\operatorname{int}\operatorname{dom} f$. Since B is compact, and so is its image under ∇g by continuity of ∇g, we obtain that both ∇g and $\nabla^2 f$ are Lipschitz on required sets, and so H_f is also Lipschitz within B, with some constant L, i.e.,

$$\|H_f(\boldsymbol{\mu}') - H_f(\boldsymbol{\mu})\| \le L\|\boldsymbol{\mu}' - \boldsymbol{\mu}\| \ .$$

Since $\operatorname{range}(H_f(\boldsymbol{\mu}')) = \operatorname{range}(H_f(\boldsymbol{\mu})) = \mathcal{G}(f)$, this implies that

$$H_f(\boldsymbol{\mu}') \simeq H_f(\boldsymbol{\mu}) \pm L\|\boldsymbol{\mu}' - \boldsymbol{\mu}\| P \ ,$$

where P is the projection on $\mathcal{G}(f)$. Since $\operatorname{range}(H_f(\boldsymbol{\mu})) = \mathcal{G}(f)$, we have $P \preceq \sigma^{-1} H_f(\boldsymbol{\mu})$ where $\sigma = \lambda_{\min}(H_f(\boldsymbol{\mu}))$ is the smallest positive eigenvalue of $H_f(\boldsymbol{\mu})$. Thus, we have

$$H_f(\boldsymbol{\mu}') \simeq \big(1 \pm L\sigma^{-1}\|\boldsymbol{\mu}' - \boldsymbol{\mu}\|\big) H_f(\boldsymbol{\mu}) \ . \qquad \square$$

A.4 Lipschitz Gradients and Strong Convexity

In addition to (or instead of) convexity$^+$, some of our results require Lipschitz gradients or, dually, strong convexity. To be precise, we say that a differentiable function $f : \mathbb{R}^d \to \mathbb{R}$ has a Lipschitz gradient if there exists a constant L such that $\|\nabla f(\boldsymbol{u}) - \nabla f(\boldsymbol{v})\| \le L\|\boldsymbol{u} - \boldsymbol{v}\|$ for all $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^d$. If f is twice differentiable, it suffices to check that $\nabla^2 f(\boldsymbol{u}) \preceq L I_d$ for all $\boldsymbol{u}$, where $I_d \in \mathbb{R}^{d \times d}$ is the identity.

We say that f is strongly convex with the strong convexity constant σ if

$$f(\boldsymbol{v}) \ge f(\boldsymbol{u}) + \boldsymbol{g}^\intercal(\boldsymbol{v} - \boldsymbol{u}) + \frac{1}{2}\sigma\|\boldsymbol{v} - \boldsymbol{u}\|^2 \ ,$$

for all $\boldsymbol{v}, \boldsymbol{u} \in \mathbb{R}^d$ and $\boldsymbol{g} \in \partial f(\boldsymbol{u})$. A standard convex analysis result states that if $f : \mathbb{R}^d \to \mathbb{R}$ has a gradient with Lipschitz constant L then f^* is strongly convex with the strong convexity constant $\sigma = 1/L$ (see Prop. 12.60 of Rockafellar and Wets, 2009).

B Proofs from Section 4

B.1 Proof of Theorem 4.1

We first express the market-clearing equilibrium definition using the trader potential functions F_i instead of trader utilities. Since $[-F_i(-\boldsymbol{r}_i) + c_i]$ is a monotone one-to-one transformation of the utility $U_i(\boldsymbol{r}_i, c_i)$, we get the following equivalences

$$\begin{aligned}\mathbf{0} \in \operatorname*{argmax}_{\boldsymbol{\delta}\in\mathbb{R}^K} U_i(\bar{\boldsymbol{r}}_i + \boldsymbol{\delta},\, \bar{c}_i - \boldsymbol{\delta}\cdot\bar{\boldsymbol{\mu}}) \;\; \text{iff} \;\; \mathbf{0} &\in \operatorname*{argmin}_{\boldsymbol{\delta}\in\mathbb{R}^K} \big[F_i(-\bar{\boldsymbol{r}}_i - \boldsymbol{\delta}) - \bar{c}_i + \boldsymbol{\delta}\cdot\bar{\boldsymbol{\mu}}\big] \\ &\text{iff} \;\; \nabla F_i(-\bar{\boldsymbol{r}}_i) = \bar{\boldsymbol{\mu}} \;\;,\end{aligned}$$

where the last step follows by setting the gradient of the objective to zero at $\boldsymbol{\delta} = \mathbf{0}$. Thus, we have that $(\bar{\boldsymbol{r}}, \bar{\boldsymbol{c}}, \bar{\boldsymbol{\mu}})$ is a market-clearing equilibrium iff

$$\sum_{i=1}^N \bar{\boldsymbol{r}}_i = \mathbf{0}, \quad \sum_{i=1}^N \bar{c}_i = 0, \quad \nabla F_i(-\bar{\boldsymbol{r}}_i) = \bar{\boldsymbol{\mu}} \text{ for all } i \in [N]. \tag{20}$$

We now analyze the minimization of the potential $\sum_i F_i(-\boldsymbol{r}_i)$ subject to the market clearing constraint $\sum_i \boldsymbol{r}_i = \mathbf{0}$. We express this constraint using the convex indicator function $\mathbb{I}\{\cdot\}$, which equals zero if its argument is true and ∞ when its false. We also introduce the matrix $A \in \mathbb{R}^{K\times NK}$ with the block structure $A \coloneqq (I_K \; I_K \; \cdots \; I_K)$ where I_K is the $K \times K$ identity matrix. Thus, A implements the summation over the blocks $\boldsymbol{r}_i$, since $A\boldsymbol{r} = \sum_i \boldsymbol{r}_i$. With this notation, the potential minimization problem can be written as

$$\min_{\boldsymbol{r}\in\mathbb{R}^{NK}} \Big[\underbrace{\sum_{i=1}^N F_i(-\boldsymbol{r}_i)}_{f(-\boldsymbol{r})} + \underbrace{\mathbb{I}\{A\boldsymbol{r} = 0\}}_{g(A\boldsymbol{r})}\Big] \;\;, \tag{21}$$

where we introduced the functions $f(\boldsymbol{r}) = \sum_i F_i(\boldsymbol{r}_i)$ and $g(\boldsymbol{s}) = \mathbb{I}\{\boldsymbol{s} = \mathbf{0}\}$ for $\boldsymbol{s} \in \mathbb{R}^K$. Now, if certain conditions are satisfied, we can apply Fenchel's duality (Theorem A.4) and obtain that the value of the primal (21) equals the value of the following dual problem

$$\max_{\boldsymbol{\mu}\in\mathbb{R}^K} \Big[-f^*(A^\intercal\boldsymbol{\mu}) - g^*(\boldsymbol{\mu})\Big] \;\;,$$

which is equivalent to

$$\max_{\boldsymbol{\mu}\in\mathbb{R}^K} \Big[-\sum_{i=1}^N F_i^*(\boldsymbol{\mu})\Big] \;\;, \tag{22}$$

because $g^*(\boldsymbol{\mu}) = 0$ and $f^*(\boldsymbol{y}) = \sum_i F_i^*(\boldsymbol{y}_i)$, for $\boldsymbol{y} \in \mathbb{R}^{NK}$, so $f^*(A^\intercal\boldsymbol{\mu}) = \sum_i F_i^*(\boldsymbol{\mu})$. It remains to verify that the preconditions of Theorem A.4 are satisfied. First, we need to check that there exists $\boldsymbol{r}$ such that $A\boldsymbol{r} \in \operatorname{ri}(\operatorname{dom} g)$. Since $\operatorname{ri}(\operatorname{dom} g) = \{\mathbf{0}\}$, the vector $\boldsymbol{r} = \mathbf{0}$ satisfies this. We also need to check that there exist $\boldsymbol{\mu}$ such that $A^\intercal\boldsymbol{\mu} \in \operatorname{ri}(\operatorname{dom} f^*)$, which for our choices of A and f is equivalent to $\boldsymbol{\mu} \in \operatorname{ri}(\operatorname{dom} F_i^*)$ for all $i \in [N]$. Since $\operatorname{dom} F_i^* = \operatorname{dom} T^* = \mathcal{M}$, any $\boldsymbol{\mu} \in \mathcal{M}$ satisfies this. Thus, conclusions of Theorem A.4 hold.

The conclusions state that both the primal and the dual are attained, and $\hat{\boldsymbol{r}}$ and $\hat{\boldsymbol{\mu}}$ are their solutions if and only if $A^\intercal \hat{\boldsymbol{\mu}} = \nabla f(-\hat{\boldsymbol{r}})$ and $\hat{\boldsymbol{\mu}} \in \partial g(A\hat{\boldsymbol{r}})$. From the definitions of A and f, the first condition is equivalent to

$$\hat{\boldsymbol{\mu}} = \nabla F(-\hat{\boldsymbol{r}}_i) = \nabla T(\tilde{\boldsymbol{\theta}}_i - a_i \hat{\boldsymbol{r}}_i) \ .$$

The second condition is by conjugacy eqivalent to $A\hat{\boldsymbol{r}} = \nabla g(\hat{\boldsymbol{\mu}}) = \mathbf{0}$, i.e.,

$$\sum_{i=1}^N \hat{\boldsymbol{r}}_i = \mathbf{0} \ .$$

This establishes that $\hat{\boldsymbol{r}}$ and $\hat{\boldsymbol{\mu}}$ are solutions to the primal (21) and dual (22), if and only if they satisfy the conditions in (20), i.e., if and only if they form a market-clearing equilibrium. This proves the theorem except for the uniqueness of the equilibrium prices $\bar{\boldsymbol{\mu}}$. The uniqueness follows from the fact that $\bar{\boldsymbol{\mu}}$ minimizes $\sum_i F_i^*(\boldsymbol{\mu})$, and the functions F_i^* are strongly convex on their domain $\mathcal{M}$, which in turn follows because the functions F_i have Lipschitz gradients (a property they inherit from the log partition function T).

B.2 Proof of Theorem 4.2

We proceed similarly to the proof of Theorem 4.1. We first express the market-maker equilibrium definition using trader potentials instead of trader utilities:

$$\begin{aligned}
\mathbf{0} \in \operatorname*{argmax}_{\boldsymbol{\delta} \in \mathbb{R}^K} & U_i\Big(\boldsymbol{r}_i^\star + \boldsymbol{\delta},\, c_i^\star - C_b(\boldsymbol{s}^\star + \boldsymbol{\delta}) + C_b(\boldsymbol{s}^\star)\Big) \\
\text{iff} \quad \mathbf{0} \in \operatorname*{argmin}_{\boldsymbol{\delta} \in \mathbb{R}^K} & \big[F_i(-\boldsymbol{r}_i^\star - \boldsymbol{\delta}) - c_i^\star + C_b(\boldsymbol{s}^\star + \boldsymbol{\delta}) - C_b(\boldsymbol{s}^\star)\big] \\
\text{iff} \quad & \nabla F_i(-\boldsymbol{r}_i^\star) = \nabla C_b(\boldsymbol{s}^\star) \ .
\end{aligned}$$

Thus, we have that $(\boldsymbol{r}^\star, \boldsymbol{c}^\star, \boldsymbol{\mu}^\star)$ is a market-maker equilibrium iff, for the market state $\boldsymbol{s}^\star = \sum_{i=1}^N \boldsymbol{r}_i^\star$,

$$\sum_{i=1}^N c_i^\star + C_b(\boldsymbol{s}^\star) - C_b(\mathbf{0}) = 0, \quad \boldsymbol{\mu}^\star = \nabla C_b(\boldsymbol{s}^\star), \quad \nabla F_i(-\bar{\boldsymbol{r}}_i) = \nabla C_b(\boldsymbol{s}^\star) \text{ for all } i \in [N]. \tag{23}$$

We next use Fenchel's duality to analyze minimization of the potential $F(\boldsymbol{r}) = \sum_i F_i(-\boldsymbol{r}_i) + C_b(\sum_i \boldsymbol{r}_i)$. We again define $A \coloneqq (I_K \ I_K \ \cdots \ I_K)$ and $f(\boldsymbol{r}) \coloneqq \sum_i F_i(\boldsymbol{r}_i)$, but in this case set $g(\boldsymbol{s}) = C_b(\boldsymbol{s})$. Therefore, by Theorem A.4, we obtain the following correspondence between the primal and the dual:

$$\min_{\boldsymbol{r} \in \mathbb{R}^{NK}} \Big[\sum_{i=1}^N F_i(-\boldsymbol{r}_i) + C_b(A\boldsymbol{r})\Big] = \min_{\boldsymbol{r} \in \mathbb{R}^{NK}} \Big[f(-\boldsymbol{r}) + g(A\boldsymbol{r})\Big] \tag{24}$$

$$= \max_{\boldsymbol{\mu} \in \mathbb{R}^K} \Big[-f^*(A^\intercal \boldsymbol{\mu}) - g^*(\boldsymbol{\mu})\Big] = \max_{\boldsymbol{\mu} \in \mathbb{R}^K} \Big[-\sum_{i=1}^N F_i^*(\boldsymbol{\mu}) - bC^*(\boldsymbol{\mu})\Big] \ , \tag{25}$$

where we used the fact that $g(\boldsymbol{r}) = bC(\boldsymbol{r}/b)$ and therefore $g^*(\boldsymbol{\mu}) = bC^*(\boldsymbol{r})$ (this is immediate from the definition of conjugate). It remains to check the preconditions of Theorem A.4. First, we need

to check that there exists $\boldsymbol{r}$ such that $A\boldsymbol{r} \in \operatorname{ri}(\operatorname{dom} g)$. Since $\operatorname{ri}(\operatorname{dom} C_b) = \mathbb{R}^K$, this is vacuous and any vector $\boldsymbol{r}$ satisfies this. We also need to check that there exist $\boldsymbol{\mu}$ such that $A^\intercal \boldsymbol{\mu} \in \operatorname{ri}(\operatorname{dom} f^*)$, which is equivalent to $\boldsymbol{\mu} \in \operatorname{ri}(\operatorname{dom} F_i^*)$ for all $i \in [N]$. Since $\operatorname{dom} F_i^* = \operatorname{dom} T^* = \mathcal{M}$, any $\boldsymbol{\mu} \in \mathcal{M}$ satisfies this. Thus, conclusions of Theorem A.4 hold.

The conclusions state that both the primal and the dual are attained, and $\hat{\boldsymbol{r}}$ and $\hat{\boldsymbol{\mu}}$ are their solutions if and only if $A^\intercal \hat{\boldsymbol{\mu}} = \nabla f(-\hat{\boldsymbol{r}})$ and $\hat{\boldsymbol{\mu}} = \nabla g(A\hat{\boldsymbol{r}})$. As in the proof of Theorem 4.1, for our A and f, the first condition is equivalent to

$$\hat{\boldsymbol{\mu}} = \nabla F(-\hat{\boldsymbol{r}}_i) = \nabla T(\tilde{\boldsymbol{\theta}}_i - a_i \hat{\boldsymbol{r}}_i) \ .$$

The second condition, $g(\boldsymbol{s}) = C_b(\boldsymbol{s})$, is

$$\hat{\boldsymbol{\mu}} = \nabla C_b\big(\sum_{i=1}^N \hat{\boldsymbol{r}}_i\big) \ .$$

This establishes that $\hat{\boldsymbol{r}}$ and $\hat{\boldsymbol{\mu}}$ are solutions to the primal (24) and dual (25), if and only if they satisfy the conditions in (23), i.e., if and only if they form a market-maker equilibrium. It remains to show that $\boldsymbol{\mu}^\star$ is unique. As before this follows by strong convexity of F_i^* and the fact that $\boldsymbol{\mu}^\star$ minimizes $\sum_i F_i^*(\boldsymbol{\mu}) + bC^*(\boldsymbol{\mu})$.

C Proofs from Section 5

C.1 Proof of Theorem 5.1

By Theorem 4.1, $\bar{\boldsymbol{\mu}} = \operatorname{argmin}_{\boldsymbol{\mu} \in \mathbb{R}^K} [\sum_i F_i^*(\boldsymbol{\mu})]$ and from the first-order optimality

$$\mathbf{0} \in \partial\Big[\sum_i F_i^*(\bar{\boldsymbol{\mu}})\Big] \ .$$

Since $F_i(\boldsymbol{s}) = \frac{1}{a_i} T(\tilde{\boldsymbol{\theta}}_i + a_i \boldsymbol{s})$, the properties of the conjugates (Theorems 12.3 and 16.1 of Rockafellar, 1970) yield

$$F_i^*(\boldsymbol{\mu}) = \frac{1}{a_i}\big(T^*(\boldsymbol{\mu}) - \tilde{\boldsymbol{\theta}}_i \cdot \boldsymbol{\mu}\big) \ .$$

Thus,

$$\begin{aligned}
\sum_i F_i^*(\bar{\boldsymbol{\mu}}) &= \Big[\sum_i 1/a_i\Big] T^*(\boldsymbol{\mu}) - \Big[\sum_i \tilde{\boldsymbol{\theta}}_i/a_i\Big] \cdot \boldsymbol{\mu} \\
&\implies \partial\Big[\sum_i F_i^*(\bar{\boldsymbol{\mu}})\Big] = \Big[\sum_i 1/a_i\Big] \partial T^*(\boldsymbol{\mu}) - \Big[\sum_i \tilde{\boldsymbol{\theta}}_i/a_i\Big] \ .
\end{aligned}$$

Therefore, $\mathbf{0} \in \partial[\sum_i F_i^*(\bar{\boldsymbol{\mu}})]$ iff

$$\frac{\sum_i \tilde{\boldsymbol{\theta}}_i/a_i}{\sum_i 1/a_i} \in \partial T^*(\bar{\boldsymbol{\mu}}) \ ,$$

which is equivalent to

$$\bar{\boldsymbol{\mu}} = \nabla T\left(\frac{\sum_i \tilde{\boldsymbol{\theta}}_i/a_i}{\sum_i 1/a_i}\right) \ .$$

C.2 Proof of Theorem 5.2

We write $w_i = a_i^{-1}/(\sum_j a_j^{-1})$ for the weights in the average. Note that $N_{\text{eff}} = (\sum_i w_i^2)^{-1}$. By the fact that beliefs are independent, we have:

$$\mathbb{E}\left[\sum_{i=1}^N w_i \tilde{\boldsymbol{\theta}}_i\right] = \boldsymbol{\theta}^{\text{true}} \text{ and } \mathbb{V}\left(\sum_{i=1}^N w_i \tilde{\boldsymbol{\theta}}_i\right) \preceq N_{\text{eff}}^{-1}\sigma^2 I_K.$$

By applying the multidimensional version of Chebyshev's inequality, we therefore have

$$\Pr\left(\left\|\sum_{i=1}^N w_i \tilde{\boldsymbol{\theta}}_i - \boldsymbol{\theta}^{\text{true}}\right\| > t\right) \leq \frac{K\sigma^2}{N_{\text{eff}}\, t^2}.$$

The result then follows from Theorem 5.1, the fact that $\boldsymbol{\mu}^{\text{true}} = \nabla T(\boldsymbol{\theta}^{\text{true}})$, and the Lipschitz continuity of ∇T.

For the final claim, we have

$$N_{\text{eff}} = \frac{\left(\sum_i a_i^{-1}\right)^2}{\left(\sum_i a_i^{-2}\right)} \geq \frac{\left(\sum_i a_{\max}^{-1}\right)^2}{\left(\sum_i a_{\min}^{-2}\right)} = N(a_{\min}/a_{\max})^2.$$

D Proofs and Additional Results from Section 6

D.1 Proof of Theorem 6.5

Since gradients ∇F_i are Lipschitz, the functions F_i^* are *strongly convex* over their domain, which is $\mathcal{M}$. Therefore, their sum $G(\boldsymbol{\mu}) := \sum_i F_i^*(\boldsymbol{\mu})$ is also strongly convex on $\mathcal{M}$ with some strong convexity constant σ. Since $\bar{\boldsymbol{\mu}} \in \operatorname{ri}\mathcal{M}$, G is subdifferentiable at $\bar{\boldsymbol{\mu}}$, and since $\bar{\boldsymbol{\mu}}$ minimizes G, any element $\boldsymbol{g} \in \partial G(\bar{\boldsymbol{\mu}})$ satisfies $\boldsymbol{g}^\intercal(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}) = 0$ for all $\boldsymbol{\mu} \in \mathcal{M}$. This together with strong convexity yields the lower bound $\sum_i F_i^*(\boldsymbol{\mu}) \geq \sum_i F_i^*(\bar{\boldsymbol{\mu}}) + \frac{1}{2}\sigma\|\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}\|^2$. At the same time, $C^*(\boldsymbol{\mu})$ is bounded below by a linear function of the form $C^*(\bar{\boldsymbol{\mu}}) + \boldsymbol{u}^\intercal(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}})$, because C^* is subdifferentiable at $\bar{\boldsymbol{\mu}}$ since $\bar{\boldsymbol{\mu}} \in \operatorname{ri}\mathcal{M} \subseteq \operatorname{ri}\operatorname{dom} C^*$.

Now from the optimality of $\boldsymbol{\mu}^\star$ and the lower bounds on $\sum_i F_i^*(\boldsymbol{\mu})$ and $C^*(\boldsymbol{\mu})$, we have

$$\begin{aligned}
\sum_i F_i^*(\bar{\boldsymbol{\mu}}) + bC^*(\bar{\boldsymbol{\mu}}) &\geq \sum_i F_i^*(\boldsymbol{\mu}^\star) + bC^*(\boldsymbol{\mu}^\star) \\
&\geq \Big(\sum_i F_i^*(\bar{\boldsymbol{\mu}}) + \frac{1}{2}\sigma\|\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}\|^2\Big) + b\big(C^*(\bar{\boldsymbol{\mu}}) + \boldsymbol{u}^\intercal(\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}})\big) \\
&\implies \|\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}\|^2 \leq -\frac{2b}{\sigma}\boldsymbol{u}^\intercal(\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}) \\
&\implies \|\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}\| \leq \frac{2b}{\sigma}\|\boldsymbol{u}\| \ .
\end{aligned}$$

D.2 A Remark on Partial and Incoherent Beliefs

The proof of Theorem 6.5 crucially relies on the fact that $\operatorname{dom}(\sum_i F_i^*) = \mathcal{M} \subseteq \operatorname{dom} C^*$, i.e., that the cost function C does not force any additional constraints on $\boldsymbol{\mu}$ beyond those already represented

by the trader potentials F_i. This is natural in our setting, because trader utilities restrict the equilibrium prices to lie in the smallest set including all coherent price vectors, $\mathcal{M}$. This property of the trader utilities means that the traders would be always willing to trade if the prices were outside the set $\mathcal{M}$. If the trader utilities did not have this property, for instance, if each trader was interested in only a few securities, or their beliefs were incoherent, then this result might not hold. In such a setting, we might end up with $\bar{\boldsymbol{\mu}} \notin \operatorname{dom} C^*$. At best, we could then show that

$$\lim_{b \to 0} \boldsymbol{\mu}^\star(b; C) = \operatorname*{argmin}_{\boldsymbol{\mu} \in \operatorname{dom} C^*} \left[\sum_{i=1}^{N} F_i^*(\boldsymbol{\mu}) \right].$$

This, of course, agrees with Theorem 6.5 for our specific setting when the restriction to $\operatorname{dom} C^*$ creates no additional constraints, because $\operatorname{dom}(\sum_i F_i^*) \subseteq \operatorname{dom} C^*$.

D.3 Proof of Theorem 6.6

The proof will proceed by analyzing the Taylor expansion of the dual objective characterizing $\boldsymbol{\mu}^\star$. However, the functions F_i^* and C^* might not be differentiable in the standard sense, because their domains might have empty interiors (and only non-empty relative interiors). Fortunately, F_i and C are convex$^+$, so by Proposition A.9 there exist convex$^+$ functions G_i and R that coincide with F_i^* and C^* on $\operatorname{aff} \operatorname{dom} F_i^*$ and $\operatorname{aff} \operatorname{dom} C^*$. These functions are three times continuously differentiable, which is what we need to obtain the third order Taylor expansion.

Note that $\operatorname{dom} F_i^* = \mathcal{M}$ for all $i \in [N]$. Let $\mathcal{A}$ denote the affine hull of $\mathcal{M} = \operatorname{dom} F_i^*$. Since $\mathcal{M} \subseteq \operatorname{dom} C^*$, the dual (8) is equivalent to

$$\boldsymbol{\mu}^\star = \operatorname*{argmin}_{\boldsymbol{\mu} \in \mathcal{A}} \left[\sum_{i=1}^{N} G_i(\boldsymbol{\mu}) + bR(\boldsymbol{\mu}) \right]. \tag{26}$$

Note that $\bar{\boldsymbol{\mu}} \in \operatorname{ri} \mathcal{M}$, because by the definition $\bar{\boldsymbol{\mu}} = \nabla T(\bar{\boldsymbol{s}})$ for some $\bar{\boldsymbol{s}}$ and the gradients of T are in $\operatorname{ri} \mathcal{M}$. Thus, functions G_i and R are differentiable and have Hessians at $\bar{\boldsymbol{\mu}}$ (by convexity$^+$). We apply the Taylor expansion at $\bar{\boldsymbol{\mu}}$ to analyze the value of the objective at $\boldsymbol{\mu}^\star$. Let $G(\boldsymbol{\mu}) \coloneqq \sum_i G_i(\boldsymbol{\mu})$. By Mean Value Theorem, we have

$$\begin{aligned}
\nabla G(\boldsymbol{\mu}^\star) &= \nabla G(\bar{\boldsymbol{\mu}}) + \nabla^2 G(\boldsymbol{\mu}_G)(\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}^\star) \\
\nabla R(\boldsymbol{\mu}^\star) &= \nabla R(\bar{\boldsymbol{\mu}}) + \nabla^2 R(\boldsymbol{\mu}_R)(\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}^\star)
\end{aligned}$$

for some $\boldsymbol{\mu}_G$ and $\boldsymbol{\mu}_R$ on the line segment connecting $\boldsymbol{\mu}^\star$ with $\bar{\boldsymbol{\mu}}$. By Theorem 6.5, $\|\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}\| = O(b)$ as $b \to 0$, and thus also $\|\boldsymbol{\mu}_G - \bar{\boldsymbol{\mu}}\| = O(b)$ and $\|\boldsymbol{\mu}_R - \bar{\boldsymbol{\mu}}\| = O(b)$. By the continuity of third derivatives of G and R in the neighborhood of $\bar{\boldsymbol{\mu}}$, the Hessians of G and R are Lipschitz in some neighborhood of $\bar{\boldsymbol{\mu}}$, which means that

$$\begin{aligned}
\nabla^2 G(\boldsymbol{\mu}_G) &= \nabla^2 G(\bar{\boldsymbol{\mu}}) + \Delta_G \\
\nabla^2 R(\boldsymbol{\mu}_R) &= \nabla^2 R(\bar{\boldsymbol{\mu}}) + \Delta_R
\end{aligned}$$

where Δ_G and Δ_R are matrices with $\|\Delta_G\| = O(\|\boldsymbol{\mu}_G - \bar{\boldsymbol{\mu}}\|) = O(b)$ and $\|\Delta_R\| = O(\|\boldsymbol{\mu}_R - \bar{\boldsymbol{\mu}}\|) = O(b)$.

We next calculate Hessians $\nabla^2 G(\bar{\boldsymbol{\mu}})$ and $\nabla^2 R(\bar{\boldsymbol{\mu}})$. First, note that

$$\nabla^2 G(\boldsymbol{\mu}) = \sum_i \nabla^2 G_i(\boldsymbol{\mu}) \ ,$$

and by Proposition A.9, we have

$$\nabla^2 G_i(\boldsymbol{\mu}) = H^+_{F_i}(\boldsymbol{\mu}) = (1/a_i) H^+_T(\boldsymbol{\mu}) \ ,$$

where the last equality follows, because $H_{F_i}(\boldsymbol{\mu}) = a_i H_T(\boldsymbol{\mu})$ from the definition of F_i. Thus,

$$\nabla^2 G(\boldsymbol{\mu}) = \sum_i (1/a_i) H^+_T(\boldsymbol{\mu}) = (N/\bar{a}) H^+_T(\boldsymbol{\mu}) \ .$$

We also have

$$\nabla^2 R(\boldsymbol{\mu}) = H^+_C(\boldsymbol{\mu}) \ .$$

By Proposition A.9, the affine space $\mathcal{A}$ is parallel to $\mathcal{G}(T)$. From the optimality of $\boldsymbol{\mu}^\star$ in (26), we have $\big(\nabla G(\boldsymbol{\mu}^\star) + b\nabla R(\boldsymbol{\mu}^\star)\big) \perp \mathcal{G}(T)$. Thus, writing P for the projection on $\mathcal{G}(T)$, we obtain

$$\begin{aligned}
\mathbf{0} &= P\Big(\nabla G(\boldsymbol{\mu}^\star) + b\nabla R(\boldsymbol{\mu}^\star)\Big) \\
&= P\nabla G(\bar{\boldsymbol{\mu}}) + P\nabla^2 G(\bar{\boldsymbol{\mu}})(\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}) + \underbrace{P\Delta_G(\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}})}_{\boldsymbol{\varepsilon}_G} + bP\nabla R(\bar{\boldsymbol{\mu}}) + \underbrace{bP\big[\nabla^2 R(\bar{\boldsymbol{\mu}}) + \Delta_R\big](\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}})}_{\boldsymbol{\varepsilon}_R} \qquad (27) \\
&= \nabla^2 G(\bar{\boldsymbol{\mu}})(\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}) + bP\nabla R(\bar{\boldsymbol{\mu}}) + \boldsymbol{\varepsilon}_G + \boldsymbol{\varepsilon}_R \ . \qquad (28)
\end{aligned}$$

In Eq. (27), the terms $\boldsymbol{\varepsilon}_G$ and $\boldsymbol{\varepsilon}_R$ have norms $O(b^2)$, because $\|\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}\|$, $\|\Delta_G\|$ and $\|\Delta_R\|$ are all at most $O(b)$. In Eq. (28), we use that $P\nabla G(\bar{\boldsymbol{\mu}}) = \mathbf{0}$ by optimality of $\bar{\boldsymbol{\mu}}$. We also use that $P\nabla^2 G(\bar{\boldsymbol{\mu}}) = \nabla^2 G(\bar{\boldsymbol{\mu}})$, because $\mathrm{range}(\nabla^2 G(\bar{\boldsymbol{\mu}})) = \mathcal{G}(T)$.

Since $(\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}) \in \mathcal{G}(T)$, multiplying Eq. (28) by $[\nabla^2 G(\bar{\boldsymbol{\mu}})]^+$, we obtain

$$\begin{aligned}
\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}} + \underbrace{[\nabla^2 G(\bar{\boldsymbol{\mu}})]^+(\boldsymbol{\varepsilon}_G + \boldsymbol{\varepsilon}_R)}_{\boldsymbol{\varepsilon}} &= -b[\nabla^2 G(\bar{\boldsymbol{\mu}})]^+ P\nabla R(\bar{\boldsymbol{\mu}}) \\
&= -b(\bar{a}/N) H_T(\bar{\boldsymbol{\mu}}) P\nabla R(\bar{\boldsymbol{\mu}}) \\
&= -b(\bar{a}/N) H_T(\bar{\boldsymbol{\mu}}) \partial C^*(\bar{\boldsymbol{\mu}}) \ ,
\end{aligned}$$

where the last step follows, because $H_T(\bar{\boldsymbol{\mu}})P = H_T(\bar{\boldsymbol{\mu}})$, and by Proposition A.9, $\partial C^*(\bar{\boldsymbol{\mu}}) = \nabla R(\bar{\boldsymbol{\mu}}) + \mathcal{G}^\perp(C)$, and $\mathcal{G}^\perp(C) \subseteq \mathcal{G}^\perp(T)$. The theorem now follows by noting that $\|\boldsymbol{\varepsilon}\| = O(b^2)$.

E Proofs and Additional Results for Section 7

E.1 Relationship between the Suboptimality of Potential and the Convergence Error

In this appendix, we relate the suboptimality of the objective to several other quantities used in analysis of convergence error. We begin by defining these quantities.

Given an allocation vector $\boldsymbol{r} \in R^{NK}$, the associated market price (the gradient of the cost) will be denoted $\boldsymbol{\mu}_0(\boldsymbol{r})$ and the gradients of trader potentials will be denoted $\boldsymbol{\mu}_i(\boldsymbol{r})$:

$$\boldsymbol{\mu}_0(\boldsymbol{r}) \coloneqq \nabla C_b\big(\textstyle\sum_{i=1}^N \boldsymbol{r}_i\big) = \nabla C\big(\sum_{i=1}^N \boldsymbol{r}_i/b\big) \tag{29}$$

$$\boldsymbol{\mu}_i(\boldsymbol{r}) \coloneqq \nabla F_i(-\boldsymbol{r}_i) = \nabla T(\tilde{\boldsymbol{\theta}}_i - a_i \boldsymbol{r}_i) \qquad \text{for } i \in [N] \tag{30}$$

where T is the log partition function (see Eq. 4). As in the body of the paper, let $\boldsymbol{r}^\star$ denote an arbitrary minimizer of F and let $F^\star$ denote the minimum value of F. From Theorem 4.2, at any equilibrium allocation $\boldsymbol{r}^\star$,

$$\boldsymbol{\mu}_0(\boldsymbol{r}^\star) = \boldsymbol{\mu}^\star \qquad \text{and} \qquad \boldsymbol{\mu}_i(\boldsymbol{r}^\star) = \boldsymbol{\mu}^\star \text{ for all } i \in [N].$$

Let $\boldsymbol{r}^t$ denote the allocation vector after the tth trade, and $\boldsymbol{\mu}^t \coloneqq \boldsymbol{\mu}_0(\boldsymbol{r}^t)$ be the corresponding market price.

We next show that to bound the convergence error $\|\boldsymbol{\mu}_0(\boldsymbol{r}^t) - \boldsymbol{\mu}^\star\|$ it suffices to bound the suboptimality of the current objective value, $F(\boldsymbol{r}^t) - F^\star$. In fact, we show that the suboptimality $F(\boldsymbol{r}^t) - F^\star$ simultaneously bounds $\|\boldsymbol{\mu}_i(\boldsymbol{r}^t) - \boldsymbol{\mu}^\star\|$, which can be viewed as a measure of suboptimality of individual traders and will be used in our later analysis:

Theorem E.1. *If ∇C has the Lipschitz constant L_C then for any $\boldsymbol{r} \in \mathbb{R}^{NK}$*

$$F(\boldsymbol{r}) - F^\star \geq \frac{b}{2L_C} \cdot \|\boldsymbol{\mu}_0(\boldsymbol{r}) - \boldsymbol{\mu}^\star\|^2 + \frac{1}{2L_T} \cdot \sum_{i=1}^N \frac{1}{a_i} \|\boldsymbol{\mu}_i(\boldsymbol{r}) - \boldsymbol{\mu}^\star\|^2 \ ,$$

where L_T is the Lipschitz constant of ∇T.

Proof. First note that since ∇C and ∇T have Lipschitz constants L_C and L_T, their conjugates are strongly convex with constants $1/L_C$ and $1/L_T$. Further, by the properties of conjugates (see Theorems 12.3 and 16.1 of Rockafellar, 1970), the definitions of F_i and C_b yield

$$F_i^*(\boldsymbol{\mu}) = \frac{1}{a_i}\big(T^*(\boldsymbol{\mu}) - \tilde{\boldsymbol{\theta}}_i \cdot \boldsymbol{\mu}\big) \ , \quad C_b^*(\boldsymbol{\mu}) = bC^*(\boldsymbol{\mu}) \ ,$$

and so F_i^* and C_b^* are strongly convex, respectively, with constants $1/(a_i L_T)$ and b/L_C.

We now invoke the duality result of Theorem 4.2 to prove our theorem. Specifically, from Eqs. (24) and (25), we have

$$F(\boldsymbol{r}^\star) = -\sum_{i=1}^N F_i^*(\boldsymbol{\mu}^\star) - C_b^*(\boldsymbol{\mu}^\star) \ .$$

Therefore, for any $\boldsymbol{r}$, we have

$$\begin{aligned} F(\boldsymbol{r}) - F^\star &= \sum_{i=1}^N F_i(-\boldsymbol{r}_i) + C_b\big(\sum_i \boldsymbol{r}_i\big) + \sum_{i=1}^N F_i^*(\boldsymbol{\mu}^\star) + C_b^*(\boldsymbol{\mu}^\star) \\ &= \sum_{i=1}^N \Big[F_i(-\boldsymbol{r}_i) + F_i^*(\boldsymbol{\mu}^\star) + \boldsymbol{r}_i^\intercal \boldsymbol{\mu}^\star\Big] + \Big[C_b\big(\sum_i \boldsymbol{r}_i\big) + C_b^*(\boldsymbol{\mu}^\star) - \big(\sum_i \boldsymbol{r}_i\big)^\intercal \boldsymbol{\mu}^\star\Big]. \end{aligned} \tag{31}$$

Using conjugacy and strong convexity, we next show that the terms in the brackets can be lower-bounded by quadratic functions.

Let $\boldsymbol{s} \coloneqq \sum_i \boldsymbol{r}_i$. Since $\boldsymbol{\mu}_i(\boldsymbol{r}) = \nabla F_i(-\boldsymbol{r}_i)$ and $\boldsymbol{\mu}_0(\boldsymbol{r}) = \nabla C_b(\boldsymbol{s})$,

$$\begin{aligned}(-\boldsymbol{r}_i) \in \partial F_i^*(\boldsymbol{\mu}_i(\boldsymbol{r})) \quad &\text{and} \quad F_i(-\boldsymbol{r}_i) = -\boldsymbol{r}_i^\mathsf{T}\boldsymbol{\mu}_i(\boldsymbol{r}) - F_i^*(\boldsymbol{\mu}_i(\boldsymbol{r})) \\ \boldsymbol{s} \in \partial C_b^*(\boldsymbol{\mu}_0(\boldsymbol{r})) \quad &\text{and} \quad C_b(\boldsymbol{\mu}_0(\boldsymbol{r})) = \boldsymbol{s}^\mathsf{T}\boldsymbol{\mu}_0(\boldsymbol{r}) - C_b^*(\boldsymbol{\mu}_0(\boldsymbol{r})) \ .\end{aligned}$$

Using these identities and invoking the strong convexity of F_i and C_b, we thus obtain

$$\begin{aligned}F_i(-\boldsymbol{r}_i) + F_i^*(\boldsymbol{\mu}^\star) + \boldsymbol{r}_i^\mathsf{T}\boldsymbol{\mu}^\star &= F_i^*(\boldsymbol{\mu}^\star) - (-\boldsymbol{r}_i)^\mathsf{T}\big(\boldsymbol{\mu}^\star - \boldsymbol{\mu}_i(\boldsymbol{r})\big) - F_i^*(\boldsymbol{\mu}_i(\boldsymbol{r})) \\ &\geq \frac{1}{2a_i L_T}\|\boldsymbol{\mu}_i(\boldsymbol{r}^t) - \boldsymbol{\mu}^\star\|^2 \\ C_b(\boldsymbol{s}) + C_b^*(\boldsymbol{\mu}^\star) - \boldsymbol{s}^\mathsf{T}\boldsymbol{\mu}^\star &= C_b^*(\boldsymbol{\mu}^\star) - \boldsymbol{s}^\mathsf{T}\big(\boldsymbol{\mu}^\star - \boldsymbol{\mu}_0(\boldsymbol{r})\big) - C_b^*(\boldsymbol{\mu}_0(\boldsymbol{r})) \\ &\geq \frac{b}{2L_C}\|\boldsymbol{\mu}_0(\boldsymbol{r}) - \boldsymbol{\mu}^\star\|^2 \ .\end{aligned}$$

The theorem now follows by applying these lower bounds in Eq. (31). □

E.2 Proof of Theorem 7.2

The result follows by Theorem E.1, using the fact that $\boldsymbol{\mu}^t = \boldsymbol{\mu}_0(\boldsymbol{r}^t)$.

E.3 Local Convergence Rate of Block-Coordinate Descent

Our local convergence-rate results are based on various local properties of F, by which we mean properties that hold on some *proper level set* of F, in the sense of the following definition:

Definition E.2 (Level Set)**.** For a given $\lambda \in \mathbb{R}$, the level set $S(F, \lambda)$ of a function F is defined as the set of points $\boldsymbol{r}$ where F is at most λ:

$$S(F, \lambda) \coloneqq \{\boldsymbol{r} : \ F(\boldsymbol{r}) \leq \lambda\} \ .$$

The level set $S(F, \lambda)$ is called *proper* if $\lambda > F^\star$.

Our first result requires that the objective be locally strongly convex and smooth in the sense of the following definition:

Definition E.3 (Strong Convexity and Smoothness for Matrix Seminorms)**.** Let A and B be symmetric positive-semidefinite matrices. We say that a differentiable function F is strongly convex on a set S with respect to A and smooth on S with respect to B if

$$\frac{1}{2}\boldsymbol{\delta}^\mathsf{T}A\boldsymbol{\delta} \leq F(\boldsymbol{r} + \boldsymbol{\delta}) - F(\boldsymbol{r}) - \boldsymbol{\delta}^\mathsf{T}\nabla F(\boldsymbol{r}) \leq \frac{1}{2}\boldsymbol{\delta}^\mathsf{T}B\boldsymbol{\delta}$$

whenever $\boldsymbol{r} \in S$ and $\boldsymbol{r} + \boldsymbol{\delta} \in S$.

To state the theorem, we introduce additional notation (the same as in the body of the paper). For a matrix $M \in \mathbb{R}^{NK \times NK}$, let $M_{\alpha,\alpha}$ denote the block consisting of rows and columns in α, i.e., $M_{\alpha,\alpha} \coloneqq E_\alpha^\mathsf{T} M E_\alpha$. And let $\mathcal{D} : \mathbb{R}^{NK \times NK} \to \mathbb{R}^{NK \times NK}$ be the operation of retaining only the block diagonal of a matrix, i.e., $\mathcal{D}(M) \coloneqq \operatorname{diag}_{\alpha \in \mathcal{A}} M_{\alpha,\alpha}$. Finally, let $\boldsymbol{g}^t \coloneqq \nabla F(\boldsymbol{r}^t)$ denote the gradient of the objective in the tth iteration.

Theorem E.4. *Assume that F attains a minimum and let $S \coloneqq S(F, \lambda)$ be a proper level set which satisfies the following conditions:*

1. *F is strongly convex and smooth on S with respect to some positive-semidefinite matrices A and B such that $\mathcal{G}(F) \subseteq \operatorname{range}(A)$, $\mathcal{G}(F) \subseteq \operatorname{range}(B)$.*
2. *There exist non-negative constants $\sigma_{\text{low}} \leq \sigma_{\text{high}} \leq \infty$ and $\ell < \infty$ such that whenever some iterate $\boldsymbol{r}^{t_0}$ lies in S, then all the consecutive iterates with $t \geq t_0 + \ell$ satisfy*

$$\sigma_{\text{low}} \mathbb{E}\left[(\boldsymbol{g}^t)^\intercal A^+ \boldsymbol{g}^t \,\middle|\, \boldsymbol{r}^{t_0}\right] \leq \mathbb{E}\left[(\boldsymbol{g}^t)^\intercal \mathcal{D}(B)^+ \boldsymbol{g}^t \,\middle|\, \boldsymbol{r}^{t_0}\right] \ ,$$
$$\mathbb{E}\left[(\boldsymbol{g}^t)^\intercal \mathcal{D}(A)^+ \boldsymbol{g}^t \,\middle|\, \boldsymbol{r}^{t_0}\right] \leq \sigma_{\text{high}} \mathbb{E}\left[(\boldsymbol{g}^t)^\intercal B^+ \boldsymbol{g}^t \,\middle|\, \boldsymbol{r}^{t_0}\right] \ ,$$

   where the expectation is over the random choice of updates by the algorithm.

Then if $F^\star < F(\boldsymbol{r}^{t_0}) < \lambda$ and $t \geq t_0 + \ell$, we have

$$\max\left\{1 - \frac{\sigma_{\text{high}}}{|\mathcal{A}|} \ , \ 0\right\} \leq \frac{\mathbb{E}\left[F(\boldsymbol{r}^{t+1}) \,\middle|\, \boldsymbol{r}^{t_0}\right] - F^\star}{\mathbb{E}\left[F(\boldsymbol{r}^t) \,\middle|\, \boldsymbol{r}^{t_0}\right] - F^\star} \leq 1 - \frac{\sigma_{\text{low}}}{|\mathcal{A}|} \ .$$

Proof. The vector of partial derivatives of F within block α will be called *partial gradient* and denoted $\nabla_\alpha F(\boldsymbol{r}) \coloneqq E_\alpha^\intercal \nabla F(\boldsymbol{r})$. The update on the block α will be denoted Ψ_α:

$$\Psi_\alpha(\boldsymbol{r}) \coloneqq \operatorname*{argmin}_{\boldsymbol{r}' \in \boldsymbol{r} + \operatorname{range}(E_\alpha)} F(\boldsymbol{r}') \ . \tag{32}$$

Since F has a minimizer, its set of gradients contains a zero, and therefore its gradient space $\mathcal{G}(F)$ coincides with the span of its gradients (see Definition A.5). We also define the set $\mathcal{G}_\alpha(F) \coloneqq E_\alpha^\intercal \mathcal{G}(F)$, which then coincides with the span of partial gradients of F. Note that by Proposition E.5, proved below, any positive-semidefinite matrix M satisfies $\operatorname{range}(M) \subseteq \operatorname{range}(\mathcal{D}(M))$. Since $\mathcal{G}(F) \subseteq \operatorname{range}(A)$, we therefore obtain $\mathcal{G}_\alpha(F) \subseteq \operatorname{range}(B_{\alpha,\alpha})$ by the following reasoning:

$$\mathcal{G}_\alpha(F) = E_\alpha^\intercal \mathcal{G}(F) \subseteq E_\alpha^\intercal \operatorname{range}(A) \subseteq E_\alpha^\intercal \operatorname{range}(\mathcal{D}(A)) = \operatorname{range}(A_{\alpha,\alpha}) \ ,$$

and similarly obtain $\mathcal{G}_\alpha(F) \subseteq \operatorname{range}(B_{\alpha,\alpha})$.

Assume $F^\star < F(\boldsymbol{r}^{t_0}) < \lambda$ and let $t \geq t_0 + \ell$. Consider the bundle $\boldsymbol{r} \coloneqq \boldsymbol{r}^t$ and analyze the value of the objective at the next iterate $\boldsymbol{r}^{t+1} = \Psi_\alpha(\boldsymbol{r})$. First note that the objective does not increase during optimization, so $F(\boldsymbol{r}) < \lambda$ and in particular $\boldsymbol{r} \in S$. From Eq. (32), we then have

$$F(\Psi_\alpha(\boldsymbol{r})) = \min_{\boldsymbol{r}' \in (\boldsymbol{r} + \operatorname{range}(E_\alpha)) \cap S} F(\boldsymbol{r}') \tag{33}$$
$$\leq \min_{\boldsymbol{r}' \in (\boldsymbol{r} + \operatorname{range}(E_\alpha)) \cap S} \left(F(\boldsymbol{r}) + (\boldsymbol{r}' - \boldsymbol{r})^\intercal \nabla F(\boldsymbol{r}) + \frac{1}{2}(\boldsymbol{r}' - \boldsymbol{r})^\intercal B(\boldsymbol{r}' - \boldsymbol{r})\right) \tag{34}$$
$$= \min_{\boldsymbol{r}' \in (\boldsymbol{r} + \operatorname{range}(E_\alpha))} \left(F(\boldsymbol{r}) + (\boldsymbol{r}' - \boldsymbol{r})^\intercal \nabla F(\boldsymbol{r}) + \frac{1}{2}(\boldsymbol{r}' - \boldsymbol{r})^\intercal B(\boldsymbol{r}' - \boldsymbol{r})\right) \tag{35}$$
$$= \min_{\boldsymbol{\delta} \in \mathbb{R}^{|\alpha|}} \left(F(\boldsymbol{r}) + \boldsymbol{\delta}^\intercal \nabla_\alpha F(\boldsymbol{r}) + \frac{1}{2}\boldsymbol{\delta}^\intercal B_{\alpha,\alpha} \boldsymbol{\delta}\right) \tag{36}$$
$$= F(\boldsymbol{r}) - \frac{1}{2}\big(\nabla_\alpha F(\boldsymbol{r})\big)^\intercal B_{\alpha,\alpha}^+ \big(\nabla_\alpha F(\boldsymbol{r})\big) \ . \tag{37}$$

In Eq. (33) it suffices to consider minimization over S, because the objective does not increase during the optimization and $\boldsymbol{r} \in S$. In Eq. (34) we use the fact that F is smooth on S with respect to B. Eq. (35) follows, because the minimum in Eq. (35) is actually attained in S. We show this by contradiction. Let $F'(\boldsymbol{r}')$ denote the function minimized in Eq. (35) and assume that the minimum of F' over $\boldsymbol{r} + \mathrm{range}(E_\alpha)$ is attained at $\boldsymbol{r}' \notin S$. Since $\boldsymbol{r} \in S$, the line connecting $\boldsymbol{r}$ and $\boldsymbol{r}'$ intersects the boundary of S at some point $\boldsymbol{r}''$, where $F(\boldsymbol{r}'') = \lambda$ by continuity of F. From the foregoing, we then have

$$F'(\boldsymbol{r}') \leq F'(\boldsymbol{r}) = F(\boldsymbol{r}) < \lambda = F(\boldsymbol{r}'') \leq F'(\boldsymbol{r}'') \ .$$

This however contradicts the convexity of F' along the line connecting $\boldsymbol{r}$ and $\boldsymbol{r}'$, because $\boldsymbol{r}''$ lies between $\boldsymbol{r}$ and $\boldsymbol{r}'$, and thus we should have $F'(\boldsymbol{r}'') \leq \max\{F'(\boldsymbol{r}), F'(\boldsymbol{r}')\}$. This means that the minimum in Eq. (35) is indeed attained somewhere in S. In Eq. (36), we make the substitution $\boldsymbol{r}' - \boldsymbol{r} = E_\alpha \boldsymbol{\delta}$, and Eq. (37) then follows by Proposition A.3, because $\nabla_\alpha F(\boldsymbol{r}) \in \mathcal{G}_\alpha(F) \subseteq \mathrm{range}(B_{\alpha,\alpha})$.

Taking expectation over the uniformly random choice of the block α, we have

$$\begin{aligned} F(\boldsymbol{r}) - \mathbb{E}_\alpha \left[F(\Psi_\alpha(\boldsymbol{r}))\right] &\geq \frac{1}{2|\mathcal{A}|} \sum_{\alpha \in \mathcal{A}} \big(\nabla_\alpha F(\boldsymbol{r})\big)^\intercal B^+_{\alpha,\alpha} \big(\nabla_\alpha F(\boldsymbol{r})\big) \\ &= \frac{1}{2|\mathcal{A}|} \big(\nabla F(\boldsymbol{r})\big)^\intercal \mathcal{D}(B)^+ \big(\nabla F(\boldsymbol{r})\big) \ . \end{aligned} \tag{38}$$

We can also apply the lower bound on F and obtain

$$\begin{aligned} F(\Psi_\alpha(\boldsymbol{r})) &\geq \min_{\boldsymbol{r}' \in (\boldsymbol{r} + \mathrm{range}(E_\alpha)) \cap S} \left(F(\boldsymbol{r}) + (\boldsymbol{r}' - \boldsymbol{r})^\intercal \nabla F(\boldsymbol{r}) + \frac{1}{2} (\boldsymbol{r}' - \boldsymbol{r})^\intercal A (\boldsymbol{r}' - \boldsymbol{r}) \right) \\ &\geq \min_{\boldsymbol{r}' \in (\boldsymbol{r} + \mathrm{range}(E_\alpha))} \left(F(\boldsymbol{r}) + (\boldsymbol{r}' - \boldsymbol{r})^\intercal \nabla F(\boldsymbol{r}) + \frac{1}{2} (\boldsymbol{r}' - \boldsymbol{r})^\intercal A (\boldsymbol{r}' - \boldsymbol{r}) \right) \qquad (39) \\ &= \min_{\boldsymbol{\delta} \in \mathbb{R}^{|\alpha|}} \left(F(\boldsymbol{r}) + \boldsymbol{\delta}^\intercal \nabla_\alpha F(\boldsymbol{r}) + \frac{1}{2} \boldsymbol{\delta}^\intercal A_{\alpha,\alpha} \boldsymbol{\delta} \right) \\ &= F(\boldsymbol{r}) - \frac{1}{2} \big(\nabla_\alpha F(\boldsymbol{r})\big)^\intercal A^+_{\alpha,\alpha} \big(\nabla_\alpha F(\boldsymbol{r})\big) \ . \qquad (40) \end{aligned}$$

The steps are analogous as in analyzing the upper bound except for Eq. (39), which is now more straightforward, since the minimum over a larger set cannot lie above a minimum over a smaller set.

Taking expectation over α, we thus obtain

$$F(\boldsymbol{r}) - \mathbb{E}_\alpha \left[F(\Psi_\alpha(\boldsymbol{r}))\right] \leq \frac{1}{2|\mathcal{A}|} \big(\nabla F(\boldsymbol{r})\big)^\intercal \mathcal{D}(A)^+ \big(\nabla F(\boldsymbol{r})\big) \ . \tag{41}$$

The same reasoning that gave us bounds on $F(\Psi_\alpha(\boldsymbol{r}))$ can be also used to bound the optimal value $F^\star$ by noting that $F^\star = F(\Psi_{\alpha^\star}(\boldsymbol{r}))$ where $\alpha^\star$ is the block containing all the coordinates, i.e., $\alpha^\star = [N] \times [K]$. (Note that $\alpha^\star$ is not necessarily in $\mathcal{A}$.) Eqs. (37) and (40) thus become

$$F^\star \leq F(\boldsymbol{r}) - \frac{1}{2} \big(\nabla F(\boldsymbol{r})\big)^\intercal B^+ \big(\nabla F(\boldsymbol{r})\big) \ , \tag{42}$$

$$F^\star \geq F(\boldsymbol{r}) - \frac{1}{2} \big(\nabla F(\boldsymbol{r})\big)^\intercal A^+ \big(\nabla F(\boldsymbol{r})\big) \ . \tag{43}$$

To finish the proof, we take the conditional expectations in Eq. (38), use the definition of σ_{low} and take the conditional expectations in Eq. (43) to obtain

$$\begin{aligned}\mathbb{E}\left[F(\boldsymbol{r}^t) \,\middle|\, \boldsymbol{r}^{t_0}\right] - \mathbb{E}\left[F(\boldsymbol{r}^{t+1}) \,\middle|\, \boldsymbol{r}^{t_0}\right] &\geq \frac{1}{2|\mathcal{A}|}\mathbb{E}\left[(\boldsymbol{g}^t)^\intercal \mathcal{D}(B)^+ \boldsymbol{g}^t \,\middle|\, \boldsymbol{r}^{t_0}\right] \\ &\geq \frac{1}{2|\mathcal{A}|}\cdot \sigma_{\text{low}}\mathbb{E}\left[(\boldsymbol{g}^t)^\intercal A^+ \boldsymbol{g}^t \,\middle|\, \boldsymbol{r}^{t_0}\right] \\ &\geq \frac{\sigma_{\text{low}}}{|\mathcal{A}|}\Big(\mathbb{E}\left[F(\boldsymbol{r}^t) \,\middle|\, \boldsymbol{r}^{t_0}\right] - F^\star\Big) \ . \end{aligned} \tag{44}$$

Similarly,

$$\mathbb{E}\left[F(\boldsymbol{r}^t) \,\middle|\, \boldsymbol{r}^{t_0}\right] - \mathbb{E}\left[F(\boldsymbol{r}^{t+1}) \,\middle|\, \boldsymbol{r}^{t_0}\right] \leq \frac{\sigma_{\text{high}}}{|\mathcal{A}|}\Big(\mathbb{E}\left[F(\boldsymbol{r}^t) \,\middle|\, \boldsymbol{r}^{t_0}\right] - F^\star\Big) \ .$$

Since the objective never increases, we can actually write

$$\mathbb{E}\left[F(\boldsymbol{r}^t) \,\middle|\, \boldsymbol{r}^{t_0}\right] - \mathbb{E}\left[F(\boldsymbol{r}^{t+1}) \,\middle|\, \boldsymbol{r}^{t_0}\right] \leq \min\left\{\frac{\sigma_{\text{high}}}{|\mathcal{A}|}\ ,\ 1\right\}\cdot\Big(\mathbb{E}\left[F(\boldsymbol{r}^t) \,\middle|\, \boldsymbol{r}^{t_0}\right] - F^\star\Big) \ . \tag{45}$$

The theorem now follows by rearranging terms in Eqs. (44) and (45). □

It remains to prove the following proposition, which was used in the proof:

Proposition E.5. *For any positive-semidefinite matrix M, we have* $\operatorname{range}(M) \subseteq \operatorname{range}(\mathcal{D}(M))$.

Proof. For $\alpha \in \mathcal{A}$, let $P_\alpha \coloneqq E_\alpha E_\alpha^T$ be the projection matrix into range E_α, and note that

$$\mathcal{D}(\mathcal{M}) = \operatorname{diag}_\alpha M_{\alpha,\alpha} = \sum_\alpha P_\alpha M P_\alpha \ .$$

Let $\boldsymbol{u} \in \mathbb{R}^{NK}$ and let $x_\alpha = \|M^{1/2}P_\alpha \boldsymbol{u}\|$. Then

$$\begin{aligned}\boldsymbol{u}^\intercal M \boldsymbol{u} \ &= \ \|M^{1/2}\boldsymbol{u}\|^2 \ = \ \Big\|M^{1/2}\big(\sum_{\alpha\in\mathcal{A}} P_\alpha \boldsymbol{u}\big)\Big\|^2 \\ &\leq \ \big(\sum_{\alpha\in\mathcal{A}} \|M^{1/2}P_\alpha\boldsymbol{u}\|\big)^2 \ = \ \big(\sum_{\alpha\in\mathcal{A}} x_\alpha\big)^2 \\ &\leq \ |\mathcal{A}|\cdot\sum_{\alpha\in\mathcal{A}} x_\alpha^2 \qquad (46) \\ &= \ |\mathcal{A}|\cdot\sum_{\alpha\in\mathcal{A}} \boldsymbol{u}^\intercal P_\alpha M P_\alpha \boldsymbol{u} \ = \ |\mathcal{A}|\cdot\big(\boldsymbol{u}^\intercal \mathcal{D}(M)\boldsymbol{u}\big) \ , \end{aligned}$$

where in Eq. (46) we used the inequality between the arithmetic and quadratic mean. Thus any $\boldsymbol{u} \in \operatorname{range}(M)$ is also in $\operatorname{range}(\mathcal{D}(M))$. □

We finish this section by showing how to apply Theorem E.4 to obtain conditional upper and lower bounds on local convergence rate, which appear in definitions in Section 7.

For the lower bounds, we assume that the optimization problem is *non-degenerate* with respect to the algorithm in the sense that the probability of reaching an iterate $\boldsymbol{r}^t$ which attains a minimum is zero. Since the optimization starts at a deterministic point $\boldsymbol{r}^0 = \boldsymbol{0}$, and the randomization is among a finite set of choices, there is only a finite set of allocation vectors $\boldsymbol{r}^t$ that can be reached at any given iteration t. Under non-degeneracy, none of these allocations (for any t) are the actual minimizers of F.

Proposition E.6. *Let $S_{\text{in}} = S(F, \lambda_{\text{in}})$ and $S_{\text{out}} = S(F, \lambda_{\text{out}})$ be level sets with $F^\star < \lambda_{\text{in}} < \lambda_{\text{out}}$, and assume that the conditions of Theorem E.4 hold for the level set S_{out} with matrices A and B, and non-negative scalars $\sigma_{\text{low}} \le \sigma_{\text{high}} \le \infty$ and $\ell < \infty$. Then for every $\boldsymbol{r}^{t_0} \in S_{\text{in}}$ there exists $c > 0$ such that for all $t \ge t_0$*

$$\mathbb{E}\left[F(\boldsymbol{r}^t) \mid \boldsymbol{r}^{t_0}\right] - F^\star \le c\gamma_{\text{high}}^{t-t_0} \ ,$$

where $\gamma_{\text{high}} = 1 - \sigma_{\text{low}}/|\mathcal{A}|$. Furthermore, if the optimization problem is non-degenerate then for every $\boldsymbol{r}^{t_0} \in S_{\text{in}}$, which has a non-zero probability of occurring, and all $t_1 \ge t_0$, there exists $c > 0$ such that for all $t \ge t_1$

$$\mathbb{E}\left[F(\boldsymbol{r}^t) \mid \boldsymbol{r}^{t_1}\right] - F^\star \ge c\gamma_{\text{low}}^{t-t_1} \ ,$$

where $\gamma_{\text{low}} = \max\{1 - \sigma_{\text{high}}/|\mathcal{A}|, 0\}$.

Proof. For the upper bound, if $F(\boldsymbol{r}^{t_0}) = F^\star$, the bound holds for any $c > 0$. Otherwise, set $c \coloneqq \gamma_{\text{high}}^{-\ell}\big(F(\boldsymbol{r}^{t_0}) - F^\star\big)$. Since the optimization does not increase the objective, this guarantees that the bound holds for $t_0 \le t \le t_0 + \ell$. For $t \ge t_0 + \ell$, Theorem E.4 gives

$$\mathbb{E}\left[F(\boldsymbol{r}^{t+1}) \mid \boldsymbol{r}^{t_0}\right] - F^\star \le \gamma_{\text{high}} \cdot \left(\mathbb{E}\left[F(\boldsymbol{r}^t) \mid \boldsymbol{r}^{t_0}\right] - F^\star\right) \ ,$$

and the upper bound now follows by induction.

For the lower bound, note that the non-degeneracy guarantees that $F(\boldsymbol{r}^t) > F^\star$ for all $t \ge t_0$. If $\gamma_{\text{low}} = 0$, the bound holds for any c. Next consider the case when $\gamma_{\text{low}} > 0$. Pick $t_1 \ge t_0$ and set

$$c \coloneqq \min_{t_1 \le t \le t_1 + \ell} \gamma_{\text{low}}^{-\ell}\left(\mathbb{E}\left[F(\boldsymbol{r}^t) \mid \boldsymbol{r}^{t_1}\right] - F^\star\right) \ ,$$

which is non-zero, because $\gamma_{\text{low}} \in (0, 1]$ and $F(\boldsymbol{r}^t) > F^\star$. This guarantees that the bound holds up to $t = t_1 + \ell$. For larger t, it follows by Theorem E.4 and induction. □

E.4 Supporting Propositions for Theorem 7.5

Throughout the propositions below, let F be a convex$^+$ function attaining a minimum. Let P be a projection on $\mathcal{G}(F)$. Since F attains a minimum, its gradient set includes zero, and therefore in Proposition A.7 we have $\boldsymbol{a} = \boldsymbol{0}$. This means that

$$F(\boldsymbol{r}) = F(P\boldsymbol{r}) \ ,$$

so we can assume, without loss of generality, that $\boldsymbol{r}^\star \in \mathcal{G}(F)$. Let $H^\star \coloneqq \nabla^2 F(\boldsymbol{r}^\star)$. Finally, recall that $S(F, \lambda)$ denotes a level set (see Definition E.2).

Proposition E.7. *For any $\lambda > F^\star$, $S(F, \lambda) = S_0 + \mathcal{G}(F)^\perp$ where $S_0 \subseteq \mathcal{G}(F)$ is compact.*

Proof. Since $F(\boldsymbol{r}) = F(P\boldsymbol{r})$, any level set S can be written as $S = S_0 + \mathcal{G}(F)^\perp$, where $S_0 = S \cap \mathcal{G}(F)$. Now if $\lambda > F^\star$ then S is non-empty and closed and hence so is S_0. It remains to argue that it is bounded. By convexity$^+$, F is strictly convex on $\mathcal{G}(F)$, so the minimum of F on the sphere $\{\boldsymbol{r} \in \mathcal{G}(F) : \|\boldsymbol{r} - \boldsymbol{r}^\star\| = 1\}$ must be some $\lambda_1 > F^\star$. By convexity, $F(\boldsymbol{r}) - F^\star \ge \lambda_1 \|\boldsymbol{r} - \boldsymbol{r}^\star\|$ for all $\boldsymbol{r} \in \mathcal{G}(F)$. Since $S_0 \subseteq \mathcal{G}(F)$ and $F(\boldsymbol{r}) \le \lambda$ for $\boldsymbol{r} \in S_0$, the set S_0 must be bounded. □

Proposition E.8. *Let $\varepsilon_{\max} > 0$. Then there exists a constant $c > 0$ such that for all $0 < \varepsilon \le \varepsilon_{\max}$, and all $\boldsymbol{r} \in S(F, F^\star + \varepsilon^2)$, we have*

$$\|P\boldsymbol{r} - \boldsymbol{r}^\star\| \le c\varepsilon \ \ , \nabla^2 F(\boldsymbol{r}) \simeq (1 \pm c\varepsilon) H^\star \ .$$

Proof. Let $S \coloneqq S(F, F^\star + \varepsilon_{\max}^2)$ and $S_0 \coloneqq S \cap \mathcal{G}(F)$, which is compact by Proposition E.7. Let

$$\sigma_1 \coloneqq \min_{\boldsymbol{r} \in S_0} \lambda_{\min}(\nabla^2 F(\boldsymbol{r}), P) \ .$$

Note that $\lambda_{\min}(\nabla^2 F(\boldsymbol{r}), P) > 0$ on the compact set S_0, and $\nabla^2 F(\cdot)$ and $\lambda_{\min}(\cdot, P)$ are continuous, so $\sigma_1 > 0$. Therefore, F is strictly convex on S with the strict convexity constant σ_1. Using the fact that $\nabla F(\boldsymbol{r}^\star) = \boldsymbol{0}$, we then have for any $\boldsymbol{r} \in S$,

$$F(\boldsymbol{r}) = F(P\boldsymbol{r}) \geq F^\star + \frac{1}{2}\sigma_1 \|P\boldsymbol{r} - \boldsymbol{r}^\star\|^2 \ .$$

Therefore, if $\boldsymbol{r} \in S(F, F^\star + \varepsilon^2) \subseteq S$ then

$$\|P\boldsymbol{r} - \boldsymbol{r}^\star\| \leq \varepsilon\sqrt{2/\sigma_1} \ .$$

For the bound on the Hessian, note that since the third derivative of F is continuous, it is upper bounded on S_0. Therefore, the Hessian is Lipschitz with some constant L on S_0, and so $\|\nabla^2 F(\boldsymbol{r}) - H^\star\| \leq L\|P\boldsymbol{r} - \boldsymbol{r}^\star\|$. Thus, since $\operatorname{range}(\nabla^2 F(\boldsymbol{r}) - H^\star) \subseteq \mathcal{G}(F)$, we have

$$\nabla^2 F(\boldsymbol{r}) \simeq H^\star \pm L\|P\boldsymbol{r} - \boldsymbol{r}^\star\| P \ .$$

Since $\sigma_2 P \preceq H^\star$ for some $\sigma_2 > 0$, we obtain

$$\nabla^2 F(\boldsymbol{r}) \simeq (1 \pm L\sigma_2^{-1}\|P\boldsymbol{r} - \boldsymbol{r}^\star\|) H^\star \ .$$

Thus, for $\boldsymbol{r} \in S(F, F^\star + \varepsilon^2) \subseteq S$ we have

$$\nabla^2 F(\boldsymbol{r}) \simeq (1 \pm \varepsilon L\sigma_2^{-1}\sqrt{2/\sigma_1}) H^\star \ .$$

Setting $c \coloneqq \max\{\sqrt{2/\sigma_1},\, L\sigma_2^{-1}\sqrt{2/\sigma_1}\}$ then proves the proposition. □

Proposition E.9. *Let S be a convex set and A and B positive-semidefinite matrices such that $A \preceq \nabla^2 F(\boldsymbol{r}) \preceq B$ for all $\boldsymbol{r} \in S$. Then F is strongly convex and smooth on S with respect to A and B.*

Proof. Let $\boldsymbol{r} \in S$, $\boldsymbol{r} + \boldsymbol{\delta} \in S$. Then from the 2nd-order Taylor expansion, we have

$$F(\boldsymbol{r} + \boldsymbol{\delta}) - F(\boldsymbol{r}) - \boldsymbol{\delta}\nabla F(\boldsymbol{r}) = \frac{1}{2}\boldsymbol{\delta}^\intercal[\nabla^2 F(\boldsymbol{r}')]\boldsymbol{\delta} \ ,$$

where $\boldsymbol{r}' \in S$. The proposition now follows, because $A \preceq [\nabla^2 F(\boldsymbol{r}')] \preceq B$. □

Proposition E.10. *For any $\lambda > F^\star$, with probability 1, the randomized block-coordinate descent algorithm with the objective F reaches an iteration t in which $\boldsymbol{r}^t \in S(F, \lambda)$.*

Proof. Since the proposition holds for $\lambda \geq F(\boldsymbol{r}^0)$, consider the case $F^\star < \lambda < F(\boldsymbol{r}^0)$, and in particular assume $F^\star < F(\boldsymbol{r}^0)$. We prove the proposition by applying Theorem E.4.

Let $S \coloneqq S(F, F(\boldsymbol{r}^0) + 1)$ and $S_0 \coloneqq S \cap \mathcal{G}(F)$, which is compact by Proposition E.7. Let

$$c_{\min} \coloneqq \min_{\boldsymbol{r} \in S_0} \lambda_{\min}(\nabla^2 F(\boldsymbol{r}), P) \ , \quad c_{\max} \coloneqq \max_{\boldsymbol{r} \in S_0} \lambda_{\max}(\nabla^2 F(\boldsymbol{r}), P) \ .$$

Note that $\lambda_{\min}(\nabla^2 F(\boldsymbol{r}), P) > 0$ on the compact set S_0, and $\nabla^2 F(\cdot)$ and $\lambda_{\min}(\cdot, P)$ are continuous, so $c_{\min} > 0$. Similarly, since $\lambda_{\max}(\nabla^2 F(\boldsymbol{r}), P) < \infty$ on S_0, the continuity yields $c_{\max} < \infty$. Since $\nabla^2 F(\boldsymbol{r}) = \nabla^2 F(P\boldsymbol{r})$, we have that for all $\boldsymbol{r} \in S$

$$c_{\min} P \preceq \nabla^2 F(\boldsymbol{r}) \preceq c_{\max} P \ .$$

Therefore, by Proposition E.9, F is strongly convex and smooth on S with respect to $A \coloneqq c_{\min} P$ and $B \coloneqq c_{\max} P$. Let

$$\sigma_{\text{low}} \coloneqq \lambda_{\min}\big(\mathcal{D}(B)^+,\, A^+\big) = \lambda_{\min}\big(c_{\max}^{-1}\mathcal{D}(P)^+,\, c_{\min}^{-1} P\big) = \frac{c_{\min}}{c_{\max}} \cdot \lambda_{\min}(\mathcal{D}(P)^+,\, P) \ ,$$

which is positive, because $\text{range}(P) \subseteq \text{range}(\mathcal{D}(P))$ by Proposition E.5.

Now, by Proposition E.6, with $S_{\text{in}} = S(F, F(\boldsymbol{r}^0))$, $S_{\text{out}} = S(F, F(\boldsymbol{r}^0) + 1) = S$, and the above matrices A and B, the scalar σ_{low}, and $\ell = 0$, we obtain that for some constant c and $\gamma \coloneqq (1 - \sigma_{\text{low}}/|\mathcal{A}|) < 1$,

$$\mathbb{E}\left[F(\boldsymbol{r}^t)\right] - F^\star = \mathbb{E}\left[F(\boldsymbol{r}^t) \mid \boldsymbol{r}^0\right] - F^\star \le c\gamma^t \ .$$

To finish the proof, we will appeal to Borel-Cantelli lemma and show that the algorithm must reach $S(F, \lambda)$ with probability 1. Specifically, note that by the Markov inequality

$$\begin{aligned}\sum_{t=1}^{\infty} \mathbb{P}\left\{F(\boldsymbol{r}^t) \ge \lambda\right\} &= \sum_{t=1}^{\infty} \mathbb{P}\left\{F(\boldsymbol{r}^t) - F^\star \ge \lambda - F^\star\right\} \\ &\le \sum_{t=1}^{\infty} \frac{c\gamma^t}{\lambda - F^\star} = \frac{c\gamma}{(1-\gamma)(\lambda - F^\star)} < \infty \ ,\end{aligned}$$

so with probability 1, only a finite number of the events $\{F(\boldsymbol{r}^t) \ge \lambda\}$ will occur; in other words, the level set $S(F, \lambda)$ is reached with probability 1. □

E.5 Proof of Theorem 7.5

Let $\varepsilon_{\max} = \sqrt{\lambda - F^\star}$. For $n = 1, 2, \ldots$ define the level sets $S(n) \coloneqq S(F, F^\star + \varepsilon_{\max}^2/n^2)$. Thus, by Proposition E.8, for some $c > 0$, we have

$$\nabla^2 F(\boldsymbol{r}) \simeq (1 \pm \underbrace{c\varepsilon_{\max}}_{c_1}/n) H^\star \quad \text{for all } \boldsymbol{r} \in S(n) \ ,$$

where we have introduced the notation $c_1 \coloneqq c\varepsilon_{\max}$. By Proposition E.9, if $n > c_1$, then the function F is strongly convex and smooth on $S(n)$ with respect to $A(n) \coloneqq (1 - c_1/n)H^\star$ and $B(n) \coloneqq (1 + c_1/n)H^\star$. Thus, $S(n)$, $A(n)$, and $B(n)$ satisfy condition (1) of Theorem E.4 for $n > c_1$.

In the remainder, we only consider the level sets $S(n)$ for $n > c_1$. For each of them define

$$\sigma_{\text{low}}(n) \coloneqq \frac{n - c_1}{n + c_1}\sigma_{\text{low}} \ , \quad \sigma_{\text{high}}(n) \coloneqq \frac{n + c_1}{n - c_1}\sigma_{\text{high}} \ .$$

We next argue that they satisfy condition (2) of Theorem E.4. It suffices to verify that condition (2) of Theorem E.4 holds for $t = t_0 + \ell$; let's call this limited variant condition (2′). If (2′) is satisfied then also (2) is satisfied, because if $\boldsymbol{r}^{t_0} \in S(n)$, then also $\boldsymbol{r}^{t_0+k} \in S(n)$ for any $k \ge 1$, and

so we can apply condition (2′) at $t = t_0 + k + \ell$ and by taking the conditional expectation we obtain the original condition (2).

To prove that condition (2) holds for $t = t_0 + \ell$, note that $S(n) \subseteq S(1) = S(F, \lambda)$, so Eq. (10) holds whenever $\boldsymbol{r}^{t_0} = \boldsymbol{r}^{t-\ell} \in S(n)$. So, assuming that $n > c_1$ and $\boldsymbol{r}^{t_0} \in S(n)$, we obtain

$$\begin{aligned}
\sigma_{\text{low}}(n) &= \frac{n - c_1}{n + c_1}\sigma_{\text{low}} \\
&\leq \frac{n/(n+c_1)}{n/(n-c_1)} \cdot \frac{\mathbb{E}\left[\, (\boldsymbol{g}^t)^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{g}^t \;\middle|\; \boldsymbol{r}^{t_0} \,\right]}{\mathbb{E}\left[\, (\boldsymbol{g}^t)^\intercal (H^\star)^+ \boldsymbol{g}^t \;\middle|\; \boldsymbol{r}^{t_0} \,\right]} = \frac{\mathbb{E}\left[\, (\boldsymbol{g}^t)^\intercal \mathcal{D}(B(n))^+ \boldsymbol{g}^t \;\middle|\; \boldsymbol{r}^{t_0} \,\right]}{\mathbb{E}\left[\, (\boldsymbol{g}^t)^\intercal A(n)^+ \boldsymbol{g}^t \;\middle|\; \boldsymbol{r}^{t_0} \,\right]} \\
\sigma_{\text{high}}(n) &= \frac{n + c_1}{n - c_1}\sigma_{\text{high}} \\
&\leq \frac{n/(n-c_1)}{n/(n+c_1)} \cdot \frac{\mathbb{E}\left[\, (\boldsymbol{g}^t)^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{g}^t \;\middle|\; \boldsymbol{r}^{t_0} \,\right]}{\mathbb{E}\left[\, (\boldsymbol{g}^t)^\intercal (H^\star)^+ \boldsymbol{g}^t \;\middle|\; \boldsymbol{r}^{t_0} \,\right]} = \frac{\mathbb{E}\left[\, (\boldsymbol{g}^t)^\intercal \mathcal{D}(A(n))^+ \boldsymbol{g}^t \;\middle|\; \boldsymbol{r}^{t_0} \,\right]}{\mathbb{E}\left[\, (\boldsymbol{g}^t)^\intercal B(n)^+ \boldsymbol{g}^t \;\middle|\; \boldsymbol{r}^{t_0} \,\right]} \;.
\end{aligned}$$

Now invoking Proposition E.6 with $S_{\text{in}} = S(n+1)$ and $S_{\text{out}} = S(n)$, we obtain that if $\boldsymbol{r}^{t_0} \in S(n+1)$ then there exists $c > 0$ such that for all $t \geq t_0$

$$\mathbb{E}\left[F(\boldsymbol{r}^t) \mid \boldsymbol{r}^{t_0}\right] - F^\star \leq c\gamma_{\text{high}}(n)^{t-t_0} \;,$$

where $\gamma_{\text{high}}(n) \coloneqq 1 - \sigma_{\text{low}}(n)/|\mathcal{A}|$. Proposition E.10 now implies that $S(n+1)$ is reached with probability 1, so $\gamma_{\text{high}}(n)$ is a valid upper bound on the local convergence rate.

Since the optimization starts at a deterministic point $\boldsymbol{r}^0 = \mathbf{0}$, and the randomization is among a finite set of choices, there is only a finite set of allocation vectors $\boldsymbol{r}^t$ that can be reached at any given iteration t. If the optimization problem is non-degenerate, then none of these allocations (for any t) are the actual minimizers of F. In that case, Proposition E.6 yields that if $\boldsymbol{r}^{t_0} \in S(n+1)$, and it has a non-zero probability of occurring, then for all $t_1 \geq t_0$, there exists $c > 0$ such that for all $t \geq t_1$

$$\mathbb{E}\left[F(\boldsymbol{r}^t) \mid \boldsymbol{r}^{t_1}\right] - F^\star \geq c\gamma_{\text{low}}(n)^{t-t_1} \;,$$

where $\gamma_{\text{low}}(n) \coloneqq 1 - \sigma_{\text{high}}(n)/|\mathcal{A}|$. Since $S(n+1)$ is reached with probability 1, this means that any valid upper bound must be greater than $\gamma_{\text{low}}(n)$.

The theorem now follows, because $\gamma_{\text{high}}(n) \to 1 - \sigma_{\text{low}}/|\mathcal{A}|$, and $\gamma_{\text{low}} \to 1 - \sigma_{\text{high}}/|\mathcal{A}|$, as $n \to \infty$.

E.6 Local Convergence of the Market

Throughout this section, we assume that C is convex$^+$, which implies that F is convex$^+$ as well. In order to apply Theorem 7.5, we need to analyze the gradient and Hessian of F. We begin the analysis by deriving explicit expressions for ∇F and $\nabla^2 F$ using the gradients and Hessians of T and C. It will be convenient to do so for trader-level blocks ∇_i and ∇^2_{ij}.

The gradient of F is composed of blocks

$$\begin{aligned}
\nabla_i F(\boldsymbol{r}) &= -\nabla F_i(-\boldsymbol{r}_i) + \nabla C_b\big(\textstyle\sum_i \boldsymbol{r}_i\big) \\
&= -\boldsymbol{\mu}_i(\boldsymbol{r}) + \boldsymbol{\mu}_0(\boldsymbol{r}) \;,
\end{aligned} \tag{47}$$

where we plugged in the definitions of $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_0$ from Eqs. (29) and (30).

For the Hessian, first consider $\nabla_{ii} F$:

$$\nabla^2_{ii} F(\boldsymbol{r}) = \nabla^2 F_i(-\boldsymbol{r}_i) + \nabla^2 C_b\big(\textstyle\sum_i \boldsymbol{r}_i\big)$$

$$= a_i \nabla^2 T(\tilde{\boldsymbol{\theta}}_i - a_i \boldsymbol{r}_i) + \frac{1}{b} \nabla^2 C\Big((\textstyle\sum_i \boldsymbol{r}_i)/b \Big)$$
$$= a_i H_T(\boldsymbol{\mu}_i(\boldsymbol{r})) + \frac{1}{b} H_C(\boldsymbol{\mu}_0(\boldsymbol{r})) \ . \tag{48}$$

Here, recall that for a convex$^+$ function f, its Hessian at any given point is only a function of the gradient at that point, which is denoted H_f. In Eq. (48), we express $\nabla^2 T$ and $\nabla^2 C$ using the respective functions H_T and H_C and the definitions of $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_0$.

For $i \neq j$, the block $\nabla^2_{ij} F$ is

$$\nabla^2_{ij} F(\boldsymbol{r}) = \nabla^2 C_b\big(\textstyle\sum_i \boldsymbol{r}_i\big) = \frac{1}{b} \nabla^2 C\Big((\textstyle\sum_i \boldsymbol{r}_i)/b \Big)$$
$$= \frac{1}{b} H_C(\boldsymbol{\mu}_0(\boldsymbol{r})) \ . \tag{49}$$

At any optimum $\boldsymbol{r}^\star$, we have $\boldsymbol{\mu}_i(\boldsymbol{r}^\star) = \boldsymbol{\mu}_0(\boldsymbol{r}^\star) = \boldsymbol{\mu}^\star$. Thus, using the Kronecker product notation, the Hessian of F at $\boldsymbol{r}^\star$ can be expressed as

$$\nabla^2 F(\boldsymbol{r}^\star) = D \otimes H_T(\boldsymbol{\mu}^\star) + \frac{\mathbf{1}\mathbf{1}^\intercal}{b} \otimes H_C(\boldsymbol{\mu}^\star) \ , \tag{50}$$

where $D \coloneqq \operatorname{diag}_{i \in [N]} a_i$ and $\mathbf{1}$ is the N-dimensional all-ones vector

Using local Lipschitz property of H_T and H_C (Proposition A.10) and the fact that $\|\boldsymbol{\mu}^\star - \bar{\boldsymbol{\mu}}\| = O(b)$ (Theorem 6.5), we immediately obtain the following asymptotic expression for $\nabla^2 F(\boldsymbol{r}^\star)$ as $b \to 0$:

Proposition E.11. *Let $H^\star \coloneqq \nabla^2 F(\boldsymbol{r}^\star)$ and $D \coloneqq \operatorname{diag}_i a_i$. Then*

$$H^\star \simeq (1 \pm O(b)) \left(\ D \otimes H_T(\bar{\boldsymbol{\mu}}) + (\mathbf{1}\mathbf{1}^\intercal) \otimes \frac{1}{b} H_C(\bar{\boldsymbol{\mu}}) \ \right) \ .$$

We next derive asymptotic formulas for matrices $\mathcal{D}(H^\star)^+$ and $(H^\star)^+$, from which we will immediately obtain a lower bound on strong convexity via Eq. (11).

Recall that $\mathcal{A}$ is the decomposition of the coordinates $[N] \times [K]$, but for our two dynamics (`ASD` and `SSD`), this decomposition has a special structure. This structure is described by a decomposition $\mathcal{B}$ of $[K]$, which is then applied to each trader, that is $\mathcal{A} = \big\{ \{i\} \times \beta \ : \ i \in [N], \beta \in \mathcal{B} \big\}$. For $M \in R^{K \times K}$, we use the notation $\mathcal{D}_{\mathcal{B}}(M)$ to describe $\operatorname{diag}_{\beta \in \mathcal{B}} M_{\beta\beta}$. For $M \in R^{NK \times NK}$, we continue writing $\mathcal{D}(M)$ instead of a more explicit $\mathcal{D}_{\mathcal{A}}(M)$.

In stating our results, we use the following shorthands, some of which have been already introduced:

$$H^\star \coloneqq \nabla^2 F(\boldsymbol{r}^\star), \quad H_T \coloneqq H_T(\bar{\boldsymbol{\mu}}), \quad H_C \coloneqq H_C(\bar{\boldsymbol{\mu}}), \quad D \coloneqq \operatorname{diag}_i a_i, \quad P = I_N - \mathbf{1}\mathbf{1}^\intercal / N.$$

The matrix P is the projection matrix on the set of centered vectors, i.e., vectors $\boldsymbol{u}$ in $\mathbb{R}^N$ such that $\mathbf{1}^\intercal \boldsymbol{u} = 0$. With this notation, the pseudoinverses $\mathcal{D}(H^\star)^+$ and $(H^\star)^+$ are characterized in the following theorem:

Theorem E.12. *Let $M_1 \coloneqq I_N \otimes \mathcal{D}_{\mathcal{B}}(H_C)^+$ and $M_2 \coloneqq (PDP)^+ \otimes H_T^+$. Then, as $b \to 0$,*

$$\mathcal{D}(H^\star)^+ \simeq (1 \pm O(b)) \cdot b M_1 \ , \tag{51}$$

$$(H^\star)^+ \simeq M_2 \pm O(b) M_1 \ . \tag{52}$$

Local strong convexity is bounded from below by

$$\sigma_{\text{low}} = b \cdot \lambda_{\min}\Big(\ (M_2^{1/2})^+ M_1 (M_2^{1/2})^+ \ \Big) - O(b^2) \tag{53}$$

$$= b \cdot \lambda_{\min}(PDP) \cdot \lambda_{\min}\big(\ H_T^{1/2} \mathcal{D}_{\mathcal{B}}(H_C)^+ H_T^{1/2} \ \big) - O(b^2) \ , \tag{54}$$

where $\lambda_{\min}(\cdot)$ *denotes the smallest positive eigenvalue of a matrix.*

The matrices M_1 and M_2 in the statement of the theorem do not depend on the liquidity parameter b. The matrix M_2, which is the dominant part of the Hessian pseudoinverse $(H^\star)^+$, is also independent of the trader dynamics and the cost function. On the other hand, the matrix M_1 reflects the cost function and dynamics. The pseudoinverse $\mathcal{D}(H^\star)^+$ approximately equals bM_1 as $b \to 0$. The main implication is that $\sigma_{\text{low}} = \Omega(b)$. This yields linear convergence rate bound $\gamma_{\text{high}} = 1 - \Omega(b)$, which suggests worse convergence as $b \to 0$. However, in the absence of a matching lower bound, we cannot conclude that the actual convergence gets worse as $b \to 0$. We next derive a matching bound $\sigma_{\text{high}} = O(b)$ for `ASD`. Thus, for `ASD`, it is not possible to achieve a linear convergence rate better than $1 - \Theta(b)$. (We conjecture similar behavior for `SSD`.) This means there is a tradeoff between convergence, which slows down as $b \to 0$, and the market-maker bias, which gets smaller.

Proof of Theorem E.12. From Proposition E.11, we know that $H^\star \simeq (1 \pm O(b))M$ where

$$M \coloneqq D \otimes H_T + (\mathbf{1}\mathbf{1}^\intercal) \otimes \frac{1}{b} H_C \ . \tag{55}$$

Proposition E.11 also implies $\mathcal{D}(H^\star) \simeq (1 \pm O(b))\mathcal{D}(M)$. Thus, by Proposition A.1, we obtain

$$\mathcal{D}(H^\star)^+ \simeq (1 \pm O(b))\mathcal{D}(M)^+ \ , \tag{56}$$

$$(H^\star)^+ \simeq (1 \pm O(b))M^+ \ . \tag{57}$$

The analysis will therefore focus on M and convert to $H^\star$ only in the last step.

We begin with the analysis of $\mathcal{D}(M)$. Specifically, consider the block $M_{\alpha\alpha}$ where $\alpha = \{i\} \times \beta$ for some $\beta \in \mathcal{B}$. From the definition of M

$$M_{\alpha\alpha} = \frac{1}{b} H_{C,\beta\beta} + a_i H_{T,\beta\beta}.$$

Since $\text{range}(H_T) = \mathcal{G}(T) \subseteq \mathcal{G}(C) = \text{range}(H_C)$, there is some constant c_1 such that $H_T \preceq c_1 H_C$, so we can write

$$M_{\alpha\alpha} \simeq \left(\frac{1}{b} \pm a_i c_1 \right) H_{C,\beta\beta}. \tag{58}$$

Setting $c_2 = (\max_i a_i) c_1$, and combining Eq. (58) across all blocks $\alpha = \{i\} \times \beta$, we thus obtain

$$\mathcal{D}(M) \simeq \left(\frac{1}{b} \pm c_2 \right) \big(\ I_N \otimes \mathcal{D}_{\mathcal{B}}(H_C) \ \big) \ .$$

Therefore, by Proposition A.1,

$$\mathcal{D}(M)^+ \simeq \frac{b}{1 \pm b c_2} \big(\ I_N \otimes \mathcal{D}_{\mathcal{B}}(H_C)^+ \ \big)$$

$$= b(1 \pm O(b))\big(\; I_N \otimes \mathcal{D}_{\mathcal{B}}(H_C)^+ \; \big) \;\; .$$

The bound on $\mathcal{D}(H^\star)^+$ now follows by Eq. (57).

We next bound $(H^\star)^+$ by analyzing M^+. First, decompose the matrix D into blocks corresponding to the ranges of the projection matrices P and $I_N - P$. Let $A = PDP$, $B = PD(I_N - P)$ and $X = (I_N - P)D(I_N - P)$. Thus,

$$D = A + B + B^\intercal + X \;\; .$$

Using the decomposition of D, we can decompose M in order to carry out blockwise inversion:

$$\begin{aligned} M &= D \otimes H_T + \frac{N}{b}(I_N - P) \otimes H_C \\ &= A \otimes H_T + (B + B^\intercal) \otimes H_T + \left(X \otimes H_T + \frac{N}{b}(I_N - P) \otimes H_C \right) \;\; . \end{aligned} \tag{59}$$

We first analyze the Schur complement matrix, which appears in the blockwise inverse:

$$\begin{aligned} S &\coloneqq \left(X \otimes H_T + \frac{N}{b}(I_N - P) \otimes H_C \right) - (B^\intercal \otimes H_T)(A^+ \otimes H_T^+)(B^\intercal \otimes H_T) \\ &= \frac{N}{b}(I_N - P) \otimes H_C + \big(X - B^\intercal A^+ B \big) \otimes H_T \;\; . \end{aligned}$$

As we argued before, $\mathrm{range}(H_T) \subseteq \mathrm{range}(H_C)$. Also $\mathrm{range}(X) \subseteq \mathrm{range}(I_N - P)$ and $\mathrm{range}(B^\intercal A^+ B) \subseteq \mathrm{range}(I_N - P)$, so for some $c_3 > 0$, we have

$$\big(X - B^\intercal A^+ B \big) \otimes H_T \simeq \pm c_3 (I_N - P) \otimes H_C \;\; ,$$

and therefore

$$S \simeq \left(\frac{N}{b} \pm c_3 \right) \big(\; (I_N - P) \otimes H_C \; \big) \;\; . \tag{60}$$

We now apply blockwise inversion (Proposition A.2) to Eq. (59), with the bounds of Eq. (60) on the Schur complement to obtain

$$M^+ \simeq A^+ \otimes H_T^+ + \frac{b}{N \pm b c_3} Y$$

where

$$Y \coloneqq \Big(I_{NK} - (A^+ \otimes H_T^+)(B \otimes H_T) \Big) \big((I_N - P) \otimes H_C^+ \big) \Big(I_{NK} - (A^+ \otimes H_T^+)(B \otimes H_T) \Big)^\intercal$$

is a positive-semidefinite matrix. Finally, invoking Eq. (57), we obtain

$$\begin{aligned} (H^\star)^+ &\simeq (1 \pm O(b))\big(\; A^+ \otimes H_T^+ \; \big) + b(1 \pm O(b)) Y \\ &\simeq \big(\; A^+ \otimes H_T^+ \; \big) \pm O(b) \big(\; A^+ \otimes H_T^+ + Y \; \big) \\ &\simeq \big(\; A^+ \otimes H_T^+ \; \big) \pm O(b) \big(\; I_N \otimes \mathcal{D}(H_C)^+ \; \big) \;\; , \end{aligned}$$

where in the last line we used that $\mathrm{range}(A^+ \otimes H_T^+ + Y) \subseteq \mathrm{range}(I_N \otimes H_C^+) \subseteq \mathrm{range}(I_N \otimes \mathcal{D}_{\mathcal{B}}(H_C)^+)$ because $\mathrm{range}(H_T) \subseteq \mathrm{range}(H_C) \subseteq \mathrm{range}(\mathcal{D}_{\mathcal{B}}(H_C))$.

Finally, to prove Eq. (54), we use Eq. (11). First, note that $(H^\star)^+ \simeq M_2 \pm O(b)M_1$, so

$$\mathrm{range}(M_2) \subseteq \mathrm{range}(H^\star) = \mathcal{G}(F) \subseteq \mathrm{range}(I_N \otimes H_C) \subseteq \mathrm{range}(M_1) \ .$$

Therefore, if $\boldsymbol{u} \in \mathcal{G}(F) \backslash \mathrm{range}(M_2)$ we have $\boldsymbol{u}^\intercal M_1 \boldsymbol{u} > 0$, but $\boldsymbol{u}^\intercal M_2 \boldsymbol{u} = 0$, so

$$\min_{\boldsymbol{u} \in \mathcal{G}(F)\backslash\{\mathbf{0}\}} \frac{\boldsymbol{u}^\intercal M_1 \boldsymbol{u}}{\boldsymbol{u}^\intercal M_2 \boldsymbol{u}} = \min_{\boldsymbol{u} \in \mathrm{range}(M_2)\backslash\{\mathbf{0}\}} \frac{\boldsymbol{u}^\intercal M_1 \boldsymbol{u}}{\boldsymbol{u}^\intercal M_2 \boldsymbol{u}} = \lambda_{\min}(M_1, M_2) > 0$$

and so

$$\frac{\boldsymbol{u}^\intercal M_2 \boldsymbol{u}}{\boldsymbol{u}^\intercal M_1 \boldsymbol{u}} \leq \lambda_{\min}^{-1}(M_1, M_2)$$

for any $\boldsymbol{u} \in \mathcal{G}(F)\backslash\{\mathbf{0}\}$.

From the bounds in Eqs. (51) and (52), there exists a constant c such that for b sufficiently small, and for all $\boldsymbol{u} \in \mathcal{G}(F)\backslash\{\mathbf{0}\}$,

$$\begin{aligned}
\frac{\boldsymbol{u}^\intercal \mathcal{D}(H)^+ \boldsymbol{u}}{\boldsymbol{u}^\intercal H^+ \boldsymbol{u}} &\geq \frac{(1-cb)\cdot \boldsymbol{u}^\intercal (bM_1)\boldsymbol{u}}{\boldsymbol{u}^\intercal M_2 \boldsymbol{u} + cb(\boldsymbol{u}^\intercal M_1 \boldsymbol{u})} \\
&= b \cdot \frac{1-cb}{\frac{\boldsymbol{u}^\intercal M_2 \boldsymbol{u}}{\boldsymbol{u}^\intercal M_1 \boldsymbol{u}} + cb} \\
&\geq b \cdot \frac{1-cb}{\lambda_{\min}^{-1}(M_1, M_2) + cb} \\
&\geq b \cdot (\lambda_{\min}(M_1, M_2) - O(b)) \ .
\end{aligned}$$

The bound on σ_{low} now follows by Eq. (11), after noting that

$$\lambda_{\min}(M_1, M_2) = \lambda_{\min}\left((M_2^{1/2})^+ M_1 (M_2^{1/2})^+ \right) = \lambda_{\min}(PDP) \cdot \lambda_{\min}\left(H_T^{1/2} \mathcal{D}_{\mathcal{B}}(H_C)^+ H_T^{1/2} \right) \ . \qquad \square$$

E.7 Tighter Analysis of the All-securities Dynamics

In Theorem E.12 we derived a worst-case bound on the curvature, valid across all possible directions that a gradient can take. In our tighter analysis of `ASD`, we derive a tighter bound on the expected curvature, exploiting the fact that the updates are chosen uniformly at random. While our analysis only applies to `ASD`, we conjecture that a similar style of analysis can also work for `SSD`.

We will index blocks by i rather than α, since each block consists of all the coordinates controlled by the trader i.

We begin by a detailed analysis of how the block-coordinate updates affect the value of the gradient. Let $\boldsymbol{r}$ by the current iterate. Consider the update Ψ_i, which optimizes over the coordinates controlled by trader i (see Eq. 32), and let $\boldsymbol{r}' = \Psi_i(\boldsymbol{r})$ be the new iterate. By the optimality of the update, we have

$$\boldsymbol{\mu}_i(\boldsymbol{r}') = \nabla F_i(-\boldsymbol{r}_i') = \nabla C_b\big(\textstyle\sum_j \boldsymbol{r}_j'\big) = \boldsymbol{\mu}_0(\boldsymbol{r}')$$

and therefore, by Eq. (47), for all $j \in [N]$,

$$\nabla_j F(\boldsymbol{r}') = -\boldsymbol{\mu}_j(\boldsymbol{r}') + \boldsymbol{\mu}_0(\boldsymbol{r}') = -\boldsymbol{\mu}_j(\boldsymbol{r}') + \boldsymbol{\mu}_i(\boldsymbol{r}') \ .$$

Thus, after the first update, the gradient $\nabla F(\boldsymbol{r})$ can be expressed using pairwise differences of $\boldsymbol{\mu}_j(\boldsymbol{r})$. When the update Ψ_i is performed, the value of $\boldsymbol{\mu}_i$ changes, whereas $\boldsymbol{\mu}_j$ for $j \neq i$ is unchanged. The amount of change in $\boldsymbol{\mu}_i$ will be denoted as $\boldsymbol{\delta}_i$:

$$\boldsymbol{\delta}_i(\boldsymbol{r}) \coloneqq \boldsymbol{\mu}_i(\boldsymbol{r}') - \boldsymbol{\mu}_i(\boldsymbol{r}) \text{ where } \boldsymbol{r}' = \Psi_i(\boldsymbol{r}) \ .$$

This is locally bounded as follows:

Lemma E.13. *There exists constants $b_0, c > 0$ such that for every $b \leq b_0$ there exists a proper level set S such that if $\boldsymbol{r} \in S$ then*

$$\|\boldsymbol{\delta}_i(\boldsymbol{r})\| \leq cb\big\|\boldsymbol{\mu}_i(\boldsymbol{r}) - \boldsymbol{\mu}_0(\boldsymbol{r})\big\|$$

for all $i \in [N]$.

Proof. Let $\varepsilon \in (0,1)$. Since T and C are convex$^+$, by Proposition A.10 it is possible to pick δ such that

$$H_T(\boldsymbol{\mu}) \simeq (1 \pm \varepsilon) H_T(\bar{\boldsymbol{\mu}}) \ , \qquad H_C(\boldsymbol{\mu}) \simeq (1 \pm \varepsilon) H_C(\bar{\boldsymbol{\mu}})$$

whenever $\|\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}\| \leq \delta$. Pick b_0 small enough such that $\|\boldsymbol{\mu}^\star(b; C) - \bar{\boldsymbol{\mu}}\| \leq \delta/2$ for all $b \leq b_0$. Fix some $b \leq b_0$ and pick a level set S such that for all $\boldsymbol{r} \in S$, $\|\boldsymbol{\mu}_i(\boldsymbol{r}) - \boldsymbol{\mu}^\star\| \leq \delta/2$ for all i and $\|\boldsymbol{\mu}_0(\boldsymbol{r}) - \boldsymbol{\mu}^\star\| \leq \delta/2$. Thus, for any $\boldsymbol{r} \in S$, we have

$$H_T(\boldsymbol{\mu}_i(\boldsymbol{r})) \simeq (1 \pm \varepsilon) H_T(\bar{\boldsymbol{\mu}}) \text{ for all } i, \qquad H_C(\boldsymbol{\mu}_0(\boldsymbol{r})) \simeq (1 \pm \varepsilon) H_C(\bar{\boldsymbol{\mu}}) \ .$$

Now let $\boldsymbol{r} \in S$ and let $\boldsymbol{r}' = \Psi_i(\boldsymbol{r})$. By the optimality of $\boldsymbol{r}'$, we know that $\nabla_i F(\boldsymbol{r}') = \boldsymbol{0}$. From the mean value theorem applied to $\nabla_i F$, there exists some $\boldsymbol{q}$ on the line segment connecting $\boldsymbol{r}$ and $\boldsymbol{r}'$ such that

$$\boldsymbol{0} = \nabla_i F(\boldsymbol{r}') = \nabla_i F(\boldsymbol{r}) + \nabla^2_{ii} F(\boldsymbol{q})(\boldsymbol{r}' - \boldsymbol{r}) \ .$$

Since $\boldsymbol{r}'$ and $\boldsymbol{r}$ differ only in block i, we obtain

$$P(\boldsymbol{r}'_i - \boldsymbol{r}_i) = -\left(\nabla^2_{ii} F(\boldsymbol{q})\right)^+ \nabla_i F(\boldsymbol{r}) \tag{61}$$

where P is the projection on $\text{range}\big(\nabla^2_{ii} F(\boldsymbol{q})\big) = \mathcal{G}(C)$. Now applying the mean value theorem to ∇F_i, we obtain that for some $\boldsymbol{q}'$ on the line segment connecting $\boldsymbol{r}$ and $\boldsymbol{r}'$, we have

$$\begin{aligned}
\boldsymbol{\delta}_i(\boldsymbol{r}) &= \nabla F_i(-\boldsymbol{r}'_i) - \nabla F_i(-\boldsymbol{r}_i) \\
&= \nabla^2 F_i(-\boldsymbol{q}'_i)(-\boldsymbol{r}'_i + \boldsymbol{r}_i) \\
&= \nabla^2 F_i(-\boldsymbol{q}'_i) \left(\nabla^2_{ii} F(\boldsymbol{q})\right)^+ \nabla_i F(\boldsymbol{r}) \ , \qquad (62)
\end{aligned}$$

where in Eq. (62) we used Eq. (61). Now both $\boldsymbol{r}$ and $\boldsymbol{r}'$ are in the level set S and so is the line segment connecting them, which includes the points $\boldsymbol{q}$ and $\boldsymbol{q}'$. Therefore,

$$\nabla^2 F_i(-\boldsymbol{q}'_i) = a_i H_T(\boldsymbol{\mu}_i(\boldsymbol{q}')) \preceq a_i(1+\varepsilon) H_T(\bar{\boldsymbol{\mu}}) \tag{63}$$

and

$$\nabla^2_{ii} F(\boldsymbol{q}) = a_i H_T(\boldsymbol{\mu}_i(\boldsymbol{q})) + \frac{1}{b} H_C(\boldsymbol{\mu}_0(\boldsymbol{q}))$$

$$\succeq (1-\varepsilon)\left(a_i H_T(\bar{\boldsymbol{\mu}}) + \frac{1}{b} H_C(\bar{\boldsymbol{\mu}})\right)$$
$$\succeq \frac{1-\varepsilon}{b} H_C(\bar{\boldsymbol{\mu}}) \ .$$

Thus, also

$$\left(\nabla^2_{ii} F(\boldsymbol{q})\right)^+ \preceq \frac{b}{1-\varepsilon} H_C^+(\bar{\boldsymbol{\mu}}) \ . \tag{64}$$

Plugging Eq. (63) and Eq. (64) into Eq. (62), we obtain

$$\|\boldsymbol{\delta}_i(\boldsymbol{r})\| \leq a_i(1+\varepsilon)\,\|H_T(\bar{\boldsymbol{\mu}})\| \cdot \frac{b}{1-\varepsilon}\left\|H_C^+(\bar{\boldsymbol{\mu}})\right\| \|\nabla_i F(\boldsymbol{r})\| \ ,$$

finishing the proof, since $\nabla_i F(\boldsymbol{r}) = -\boldsymbol{\mu}_i(\boldsymbol{r}) + \boldsymbol{\mu}_0(\boldsymbol{r})$. □

Using the lemma, we can now prove bounds for ASD:

Theorem E.14. *Consider the all-securities dynamics. Let $M_1' \coloneqq P \otimes H_C^+$ and $M_2 \coloneqq (PDP)^+ \otimes H_T^+$. Then for every sufficiently small b, there exists a proper level set S such that*

$$\mathbb{E}\left[(\boldsymbol{g}^{t+1})^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{g}^{t+1} \,\middle|\, \boldsymbol{r}^t\right] \simeq (\boldsymbol{g}^t)^\intercal \left(\ (1 \pm O(b)) \cdot 2bM_1'\ \right)\boldsymbol{g}^t \tag{65}$$
$$\mathbb{E}\left[(\boldsymbol{g}^{t+1})^\intercal (H^\star)^+ \boldsymbol{g}^{t+1} \,\middle|\, \boldsymbol{r}^t\right] \simeq (\boldsymbol{g}^t)^\intercal \left(\ M_2 \pm O(b) M_1'\ \right)\boldsymbol{g}^t \tag{66}$$

whenever $\boldsymbol{r}^{t-1} \in S$. Local strong convexity is bounded from below and above by

$$\sigma_{\text{low}} = 2b \cdot \lambda_{\min}(PDP) \cdot \lambda_{\min}\left(\ H_T^{1/2} H_C^+ H_T^{1/2}\ \right) - O(b^2)$$
$$\sigma_{\text{high}} = 2b \cdot \lambda_{\max}(PDP) \cdot \lambda_{\max}\left(\ H_T^{1/2} H_C^+ H_T^{1/2}\ \right) + O(b^2) \ ,$$

where $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ denote the smallest and the largest positive eigenvalue of a matrix.

As mentioned above, the key consequence of Theorem E.14 is the fact that ASD converges at the rate $1 - \Theta(b)$. The key difference from Theorem E.12 is in the expression for $\mathcal{D}(H^\star)^+$. While $\mathcal{D}(H^\star)^+ \approx b(I_N \otimes H_C^+)$, as stated in Theorem E.12, when we take an expectation over an update in a single iteration, the action of $\mathcal{D}(H^\star)^+$ is equivalent to that of the matrix $2b(P \otimes H_C^+)$. Thus, the averaging effect of an expectation is to remove one rank from matrix I_N and replace it by the centering matrix $P = I_N - \mathbf{1}\mathbf{1}^\intercal/N$ (while incurring an extra factor of two). This has two consequences. First, the lower bound σ_{low} is a factor of two larger. Second, we can now obtain a non-trivial upper bound σ_{high}, which would not be possible via an analog of Eq. (11), because the range of $\mathcal{D}(H^\star)^+$ is too large. Our numerical experiments in Section 8 show that these bounds are empirically quite tight. In the body of the paper, we also show an example when the two bounds match (except for terms on the order of $O(b^2)$): when all traders have identical risk aversions and the cost function is LMSR.

We expect that the reasoning that gave rise to Theorem E.14 can be generalized to SSD, along the lines of the following conjecture:

Conjecture E.15. *For SSD, local strong convexity is bounded from below and above by*

$$\sigma_{\text{low}} = 2b \cdot \lambda_{\min}(PDP) \cdot \lambda_{\min}\left(\ H_T^{1/2} \mathcal{D}_{\mathcal{B}}(H_C)^+ H_T^{1/2}\ \right) - O(b^2)$$
$$\sigma_{\text{high}} = 2b \cdot \lambda_{\max}(PDP) \cdot \lambda_{\max}\left(\ H_T^{1/2} \mathcal{D}_{\mathcal{B}}(H_C)^+ H_T^{1/2}\ \right) + O(b^2) \ .$$

Proof of Theorem E.14. Consider b_0 and c from Lemma E.13, and let $b \le b_0$ and S be the level set from the lemma. Assume that $\boldsymbol{r}^{t-1} \in S$. After the update in the iteration $t-1$, it is guaranteed that $\boldsymbol{\mu}_0(\boldsymbol{r}^t) = \boldsymbol{\mu}_j(\boldsymbol{r}^t)$ for some j. We analyze the update in the following iteration, i.e., the iteration t. We write $\boldsymbol{r}$ for $\boldsymbol{r}^t$ and $\boldsymbol{r}'$ for $\boldsymbol{r}^{t+1}$. Let $\boldsymbol{g} \coloneqq \nabla F(\boldsymbol{r})$, $\boldsymbol{\mu}_j \coloneqq \boldsymbol{\mu}_j(\boldsymbol{r})$ for $j \in [N]$, $\boldsymbol{\mu}_0 \coloneqq \boldsymbol{\mu}_0(\boldsymbol{r})$, and similarly define $\boldsymbol{g}'$, $\boldsymbol{\mu}'_j$, $\boldsymbol{\mu}'_0$ for the iterate $\boldsymbol{r}'$.

Recall from Eq. (47) that the blocks of the gradient are

$$\boldsymbol{g}_j = \boldsymbol{\mu}_j - \boldsymbol{\mu}_0 \ .$$

Also recall that $P = I_N - \mathbf{1}\mathbf{1}^\intercal/N$. A key role in the analysis will be played by the *centered* gradient vector $\boldsymbol{u} \coloneqq (P \otimes I_K)\boldsymbol{g}$ whose blocks are

$$\boldsymbol{u}_j = \boldsymbol{\mu}_j - \hat{\boldsymbol{\mu}}$$

where $\hat{\boldsymbol{\mu}} \coloneqq (\sum_j \boldsymbol{\mu}_j)/N$ is the average among $\boldsymbol{\mu}_j$. As the final part of the setup, let $\rho = \max_j \|\boldsymbol{\mu}_j - \hat{\boldsymbol{\mu}}\|$, and since $\boldsymbol{\mu}_0 = \boldsymbol{\mu}_j$ for some $j \in [N]$, we also have $\rho \ge \|\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}\|$.

By Theorem E.12, we have

$$\mathcal{D}(H^\star)^+ = \big(1 \pm O(b)\big) b M_1 \ , \qquad (H^\star)^+ = M_2 \pm O(b) M_1 \ ,$$

where

$$\begin{aligned} M_1 &= I_N \otimes H_C^+ \\ M_2 &= (PDP)^+ \otimes H_T^+ \ . \end{aligned}$$

For the first part of the theorem (Eqs. 65 and 66), it therefore suffices to show that

$$\mathbb{E}\left[(\boldsymbol{g}')^\intercal M_1 \boldsymbol{g}' \mid \boldsymbol{r}\right] \simeq (1 \pm O(b)) \cdot \boldsymbol{g}^\intercal \big(2P \otimes H_C^+ \big) \boldsymbol{g} \tag{67}$$

$$\mathbb{E}\left[(\boldsymbol{g}')^\intercal M_2 \boldsymbol{g}' \mid \boldsymbol{r}\right] \simeq \boldsymbol{g}^\intercal M_2 \boldsymbol{g} \pm O(b) \cdot \boldsymbol{g}^\intercal \big(P \otimes H_C^+ \big) \boldsymbol{g} \ . \tag{68}$$

Note that $P \otimes H_C^+ = (P \otimes I_K) M_1 (P \otimes I_K)$ and $M_2 = (P \otimes I_K) M_2 (P \otimes I_K)$, so Eqs. (67) and (68) are equivalent to

$$\mathbb{E}\left[(\boldsymbol{g}')^\intercal M_1 \boldsymbol{g}' \mid \boldsymbol{r}\right] \simeq (1 \pm O(b)) \cdot \boldsymbol{u}^\intercal (2M_1) \boldsymbol{u} \tag{M_1-bound}$$

$$\mathbb{E}\left[(\boldsymbol{g}')^\intercal M_2 \boldsymbol{g}' \mid \boldsymbol{r}\right] \simeq \boldsymbol{u}^\intercal M_2 \boldsymbol{u} \pm O(b) \cdot \boldsymbol{u}^\intercal M_1 \boldsymbol{u} \ , \tag{M_2-bound}$$

which is what we will show next.

Assume that the ith block is chosen for an update in the iteration t and write $\boldsymbol{\delta}_i$ for $\boldsymbol{\delta}_i(\boldsymbol{r})$. Note that

$$\boldsymbol{\mu}'_0 = \boldsymbol{\mu}'_i = \boldsymbol{\mu}_i + \boldsymbol{\delta}_i \ , \qquad \boldsymbol{\mu}'_j = \begin{cases} \boldsymbol{\mu}_j & \text{if } j \ne i, \\ \boldsymbol{\mu}_i + \boldsymbol{\delta}_i & \text{if } j = i. \end{cases}$$

Thus, from Eq. (47), the blocks $\boldsymbol{g}'_j$ can be written as

$$\boldsymbol{g}'_j = \begin{cases} \boldsymbol{\mu}_i - \boldsymbol{\mu}_j + \boldsymbol{\delta}_i & \text{if } j \ne i, \\ \mathbf{0} & \text{if } j = i. \end{cases}$$

Calculate:

$$(\boldsymbol{g}')^\intercal M_1 \boldsymbol{g}' = \sum_j (\boldsymbol{g}'_j)^\intercal H_C^+ \boldsymbol{g}'_j$$

$$
\begin{aligned}
&= \sum_{j\neq i} (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j + \boldsymbol{\delta}_i)^\intercal H_C^+ (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j + \boldsymbol{\delta}_i) \\
&= \sum_{j} (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j + \boldsymbol{\delta}_i)^\intercal H_C^+ (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j + \boldsymbol{\delta}_i) - \boldsymbol{\delta}_i^\intercal H_C^+ \boldsymbol{\delta}_i \\
&= \sum_{j} \Big[(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)^\intercal H_C^+ (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) + 2\boldsymbol{\delta}_i^\intercal H_C^+ (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) + \boldsymbol{\delta}_i^\intercal H_C^+ \boldsymbol{\delta}_i \Big] - \boldsymbol{\delta}_i^\intercal H_C^+ \boldsymbol{\delta}_i \\
&\simeq \sum_{j} \Big[(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)^\intercal H_C^+ (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \Big] \pm \Big[4N \|\boldsymbol{\delta}_i\| \|H_C^+\| \rho + N \|H_C^+\| \|\boldsymbol{\delta}_i\|^2 \Big] \qquad (69) \\
&\simeq \sum_{j} \Big[(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)^\intercal H_C^+ (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \Big] \pm \Big[8Nbc \|H_C^+\| \rho^2 + 4Nb^2c^2 \|H_C^+\| \rho^2 \Big] \qquad (70) \\
&\simeq \sum_{j} \Big[(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)^\intercal H_C^+ (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \Big] \pm bc_1 \rho^2 \ . \qquad (71)
\end{aligned}
$$

In Eq. (69), we used the fact that

$$
\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\| \leq \|\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}}\| + \|\boldsymbol{\mu}_j - \hat{\boldsymbol{\mu}}\| \leq 2\rho \ .
$$

In Eq. (70), we used Lemma E.13, which implies that

$$
\|\boldsymbol{\delta}_i\| \leq bc \|\boldsymbol{\mu}_i - \boldsymbol{\mu}_0\| \leq bc \Big(\|\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}}\| + \|\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}}\| \Big) \leq 2bc\rho \ . \qquad (72)
$$

And in Eq. (70), we set $c_1 = 8Nc\|H_C^+\| + 4Nb_0c^2\|H_C^+\|$.

Recall that $\boldsymbol{u}_j = \boldsymbol{\mu}_j - \hat{\boldsymbol{\mu}}$, so $\boldsymbol{\mu}_i - \boldsymbol{\mu}_j = \boldsymbol{u}_i - \boldsymbol{u}_j$. Thus, taking expectation over i, Eq. (71) yields

$$
\begin{aligned}
\mathbb{E}_i \left[(\boldsymbol{g}')^\intercal M_1 \boldsymbol{g}' \right] &\simeq \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N \left[(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)^\intercal H_C^+ (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \right] \pm bc_1\rho^2 \\
&= \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N \left[(\boldsymbol{u}_i - \boldsymbol{u}_j)^\intercal H_C^+ (\boldsymbol{u}_i - \boldsymbol{u}_j) \right] \pm bc_1\rho^2 \\
&= \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N \left[\boldsymbol{u}_i^\intercal H_C^+ \boldsymbol{u}_i - 2\boldsymbol{u}_i^\intercal H_C^+ \boldsymbol{u}_j + \boldsymbol{u}_j H_C^+ \boldsymbol{u}_j \right] \pm bc_1\rho^2 \\
&= \sum_{i=1}^N \left[2\boldsymbol{u}_i^\intercal H_C^+ \boldsymbol{u}_i \right] \pm bc_1\rho^2 \qquad (73) \\
&= 2\boldsymbol{u}^\intercal M_1 \boldsymbol{u} \pm bc_1\rho^2 \ , \qquad (74)
\end{aligned}
$$

where Eq. (73) follows because $\sum_i \boldsymbol{u}_i = \boldsymbol{0}$. To prove ($M_1$-bound), it remains to upper bound ρ^2. Let σ be the smallest eigenvalue of H_C^+ over $\mathcal{G}(C)$; it must be greater than zero because $\text{range}(H_C) = \mathcal{G}(C)$. We can bound ρ^2 as

$$
\rho^2 \leq \sum_{i=1}^N \|\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}}\|^2 \leq \sigma^{-1} \sum_{i=1}^N (\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}})^\intercal H_C^+ (\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}}) = \sigma^{-1} \boldsymbol{u}^\intercal M_1 \boldsymbol{u} \ . \qquad (75)
$$

Plugging this bound back into Eq. (74) yields (M_1-bound).

We next prove (M_2-bound). Again, consider an update of the block i. Let $\boldsymbol{u}' \coloneqq (P \otimes I_k)\boldsymbol{g}'$ be the centered version of $\boldsymbol{g}'$. Its blocks are

$$\boldsymbol{u}'_j = \begin{cases} \boldsymbol{u}_j - \delta_i/N & \text{if } j \neq i, \\ \boldsymbol{u}_j - \delta_i/N + \delta_i & \text{if } j = i. \end{cases}$$

Thus,

$$\|\boldsymbol{u}' - \boldsymbol{u}\|^2 = (N-1) \cdot \frac{\|\delta_i\|^2}{N^2} + \left(1 - \frac{1}{N}\right)^2 \|\delta_i\|^2 \leq \|\delta_i\|^2 \leq (2bc\rho)^2$$

where the last inequality follows by Eq. (72). Also, from the definition of ρ,

$$\|\boldsymbol{u}\|^2 \leq N\rho^2 \ .$$

Now, we bound $(\boldsymbol{g}')^\intercal M_2 \boldsymbol{g}'$. Since $M_2 = (P \otimes I_K) M_2 (P \otimes I_K)$, we can write

$$\begin{aligned}
(\boldsymbol{g}')^\intercal M_2 \boldsymbol{g}' &= (\boldsymbol{u}')^\intercal M_2 \boldsymbol{u}' \\
&= \big(\boldsymbol{u} + (\boldsymbol{u}' - \boldsymbol{u})\big)^\intercal M_2 \big(\boldsymbol{u} + (\boldsymbol{u}' - \boldsymbol{u})\big) \\
&\simeq \boldsymbol{u}^\intercal M_2 \boldsymbol{u} \pm \left[2\|\boldsymbol{u}\|\|M_2\|\|\boldsymbol{u}' - \boldsymbol{u}\| + \|M_2\|\|\boldsymbol{u}' - \boldsymbol{u}\|^2\right] \\
&\simeq \boldsymbol{u}^\intercal M_2 \boldsymbol{u} \pm \left[4bc\sqrt{N}\|M_2\|\rho^2 + 4b^2c^2\|M_2\|\rho^2\right] \qquad (76) \\
&\simeq \boldsymbol{u}^\intercal M_2 \boldsymbol{u} \pm bc_2\rho^2 \ , \qquad (77)
\end{aligned}$$

where in Eq. (76) we used the previously derived bounds and in Eq. (77), we set $c_2 = 4c\sqrt{N}\|M_2\| + 4b_0c^2\|M_2\|$. Finally, using the bound on ρ^2 from Eq. (75) in Eq. (77) yields (M_2-bound).

It remains to prove the bounds σ_{low} and σ_{high}. In particular, we will show that if $\boldsymbol{r}^{t-1} \in S$ then

$$\sigma_{\text{low}} \cdot \mathbb{E}\left[(\boldsymbol{g}')^\intercal (H^\star)^+ \boldsymbol{g}' \,\middle|\, \boldsymbol{r}\right] \leq \mathbb{E}\left[(\boldsymbol{g}')^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{g}' \,\middle|\, \boldsymbol{r}\right] \leq \sigma_{\text{high}} \cdot \mathbb{E}\left[(\boldsymbol{g}')^\intercal (H^\star)^+ \boldsymbol{g}' \,\middle|\, \boldsymbol{r}\right] \qquad (78)$$

then taking expectation over $\boldsymbol{r} = \boldsymbol{r}^t$, conditionally on $\boldsymbol{r}^{t-1}$, we also obtain

$$\sigma_{\text{low}} \cdot \mathbb{E}\left[(\boldsymbol{g}')^\intercal (H^\star)^+ \boldsymbol{g}' \,\middle|\, \boldsymbol{r}^{t-1}\right] \leq \mathbb{E}\left[(\boldsymbol{g}')^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{g}' \,\middle|\, \boldsymbol{r}^{t-1}\right] \leq \sigma_{\text{high}} \cdot \mathbb{E}\left[(\boldsymbol{g}')^\intercal (H^\star)^+ \boldsymbol{g}' \,\middle|\, \boldsymbol{r}^{t-1}\right]$$

which will yield the desired conclusion by Theorem 7.5 (with the level set S and $\ell = 2$).

To prove Eq. (78), we first apply Eqs. (65) and (66):

$$\begin{aligned}
\underbrace{\frac{\mathbb{E}\left[(\boldsymbol{g}')^\intercal (H^\star)^+ \boldsymbol{g}' \,\middle|\, \boldsymbol{r}\right]}{\mathbb{E}\left[(\boldsymbol{g}')^\intercal \mathcal{D}(H^\star)^+ \boldsymbol{g}' \,\middle|\, \boldsymbol{r}\right]}}_{=:z} &\simeq (1 \pm O(b)) \cdot \frac{\boldsymbol{g}^\intercal \big(\, 2b(P \otimes I_K) M_1 (P \otimes I_K) \,\big) \boldsymbol{g}}{\boldsymbol{g}^\intercal \big(\, M_2 \ \pm \ O(b) \cdot (P \otimes I_K) M_1 (P \otimes I_K) \,\big) \boldsymbol{g}} \\
&= (1 \pm O(b)) \cdot \frac{\boldsymbol{u}^\intercal \,(\, 2bM_1 \,)\, \boldsymbol{u}}{\boldsymbol{u}^\intercal \big(\, M_2 \pm O(b) M_1 \,\big) \boldsymbol{u}} \\
&= (1 \pm O(b)) \cdot \frac{2b}{\frac{\boldsymbol{u}^\intercal M_2 \boldsymbol{u}}{\boldsymbol{u}^\intercal M_1 \boldsymbol{u}} \ \pm \ O(b)} \ .
\end{aligned}$$

Now note that the blocks of $\boldsymbol{u}$ take form $\boldsymbol{u}_i = \boldsymbol{\mu}_i - (\sum_j \boldsymbol{\mu}_j)/N$, where $\boldsymbol{\mu}_i \in \mathcal{M} \subseteq \text{range}(H_T^+)$, so $\boldsymbol{u} \in \text{range}(P \otimes H_T^+) = \text{range}(M_2) \subseteq \text{range}(M_1)$. This means that

$$0 < \lambda_{\min}(M_1, M_2) \leq \frac{\boldsymbol{u}^\intercal M_1 \boldsymbol{u}}{\boldsymbol{u}^\intercal M_2 \boldsymbol{u}} \leq \lambda_{\max}(M_1, M_2) \ ,$$

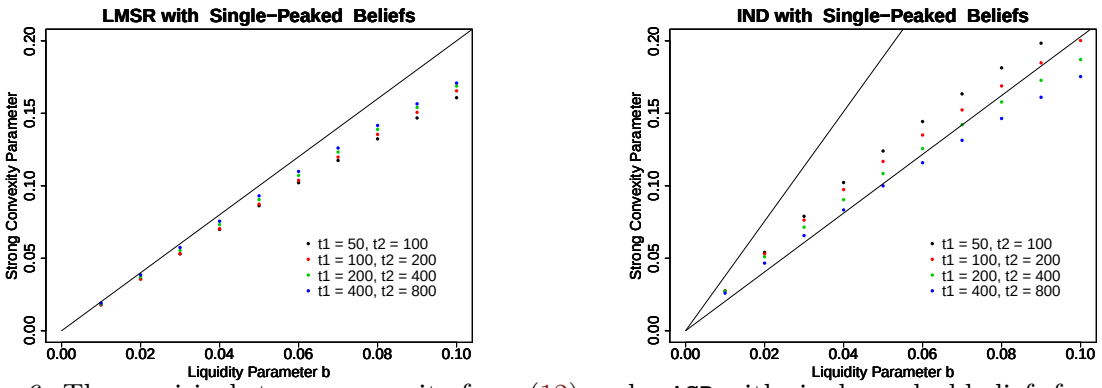

Figure 6: The empirical strong convexity from (12) under ASD with single-peaked beliefs for various values of t_1, t_2 represented as dots, asymptotic bounds for σ_{low} and σ_{high} from Theorem 7.6 (ignoring $O(b^2)$ terms) represented as lines.

so we can write

$$z \simeq (1 \pm O(b)) \cdot 2b \cdot \frac{\boldsymbol{u}^\intercal M_1 \boldsymbol{u}}{\boldsymbol{u}^\intercal M_2 \boldsymbol{u}} \ ,$$

and thus,

$$2b \cdot \lambda_{\min}(M_1, M_2) - O(b^2) \ \leq \ z \ \leq \ 2b \cdot \lambda_{\max}(M_1, M_2) + O(b^2) \ .$$

The theorem now follows, because

$$\begin{aligned} \lambda_{\min}(M_1, M_2) &= \lambda_{\min}\big(\ (M_2^+)^{1/2} M_1 (M_2^+)^{1/2} \ \big) \\ &= \lambda_{\min}(PDP) \cdot \lambda_{\min}\big(\ H_T^{1/2} H_C^+ H_T^{1/2} \ \big) \ , \end{aligned}$$

and similarly for $\lambda_{\max}(M_1, M_2)$. □

E.8 Proof of Theorem 7.6

The theorem follows from Theorems E.12 and E.14.

F Additional Results from Numerical Experiments

F.1 Tightness of the Empirical Strong Convexity for Single-Peaked Beliefs

Recall that in Figure 5, we plotted the empirical strong convexity $\hat{\sigma}$ as a function of b for all-securities dynamics and uniform beliefs. We then also present the same plot for $\hat{\sigma}$ but for single-peaked beliefs in Figure 6.

F.2 Numerical Experiments for Single-Security Dynamics

In Section 8, we provide numerical experiments evaluating the tightness of our bounds on market-maker bias and convergence error. Due to lack of space, we only reported the results for the all-securities dynamic. For completeness, we include here the analogous results for single-security dynamics. Thus, we still consider $N = 10$, $K = 5$, the risk aversion coefficients $a_i = 1$ for al $i \in [N]$, and both the uniform and single-peaked beliefs in Figure 1.

Bias/Convergence Tradeoffs. We first consider the tradeoff between bias and convergence error. Figure 7 is the analog of Figure 2 for singe-security dynamics.

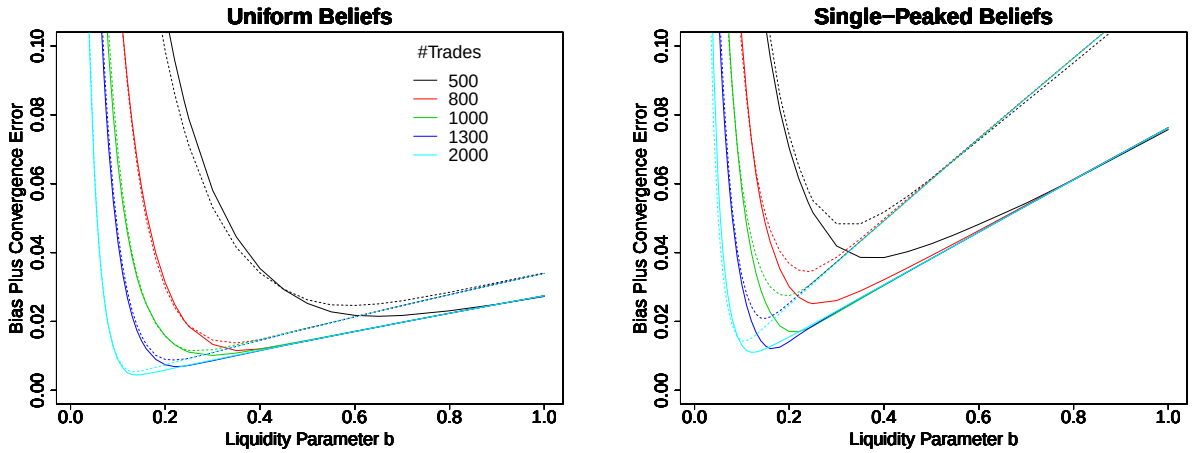

Figure 7: The tradeoff between marker-maker bias and convergence error for single-security dynamics with different beliefs and cost functions. Solid lines show the total bias and convergence error of the LMSR after various numbers of trades, averaged over 20 random trade sequences. Dotted lines show the same for IND.

Convergence Error Note that we have already considered the market-maker bias in Figure 3, which does not depend on the trader dynamics. We then turn to the convergence error alone. Again, instead of examining the convergence of prices, we examine the convergence of the objective. Figure 8 shows the empirical value of $F(\boldsymbol{r}^t) - F^\star$, averaged over 20 random sequences of trade, as a function of the number of trades, plotted on a log scale, for our two belief sets and cost functions under the single-security dynamics. As is suggested by Conjecture E.15, the plots show the convergence error for LMSR and IND right on top of each other.

When we are considering SSD, we do not have bounds for σ_{high} nor do we have a tighter analysis for σ_{low} that we have in the ASD case. However, we can compare our conjectured values for σ_{low} and σ_{high} from Conjecture E.15 with the empirical strong convexity $\hat{\sigma}$ values from (12) in Figure 9. The conjectured values are presented as dashed red lines in Figure 9. See Table 1 for a table of

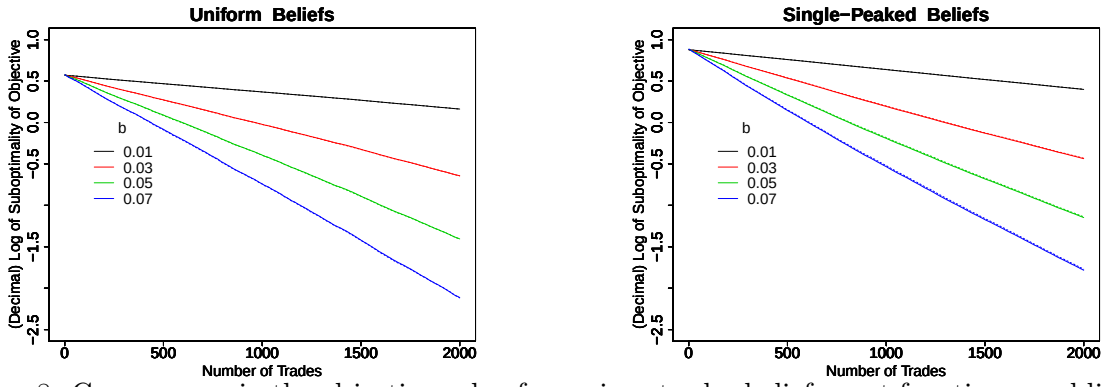

Figure 8: Convergence in the objective value for various trader beliefs, cost functions, and liquidity parameters under the single-security dynamics. Solid lines show the log error in objective for LMSR, dotted lines for IND, which are on top of each other.

Beliefs	**Dynamics**	C	σ_{low}	σ_{high}	**Conjectured** σ_{low}	**Conjectured** σ_{high}
Uniform Beliefs	ASD	LMSR	$2b$	$2b$	*	*
		IND	$2.31b$	$2.78b$	*	*
	SSD	LMSR	$1.01b$	?	$2.31b$	$2.78b$
		IND	$1.01b$	?	$2.31b$	$2.78b$
Single Peaked Beliefs	ASD	LMSR	$2b$	$2b$	*	*
		IND	$2.03b$	$3.78b$	*	*
	SSD	LMSR	$1.01b$	?	$2.03b$	$3.78b$
		IND	$1.01b$	?	$2.03b$	$3.78b$

Table 1: Table of values for the bounds σ_{low} and σ_{high} in the various cases we consider in our experiments. These were computed using Theorem 7.6 and we use the conjectured values for σ_{low} and σ_{high} from Conjecture E.15.

values for the various bounds we used in both the all-securities and single-security dynamics.

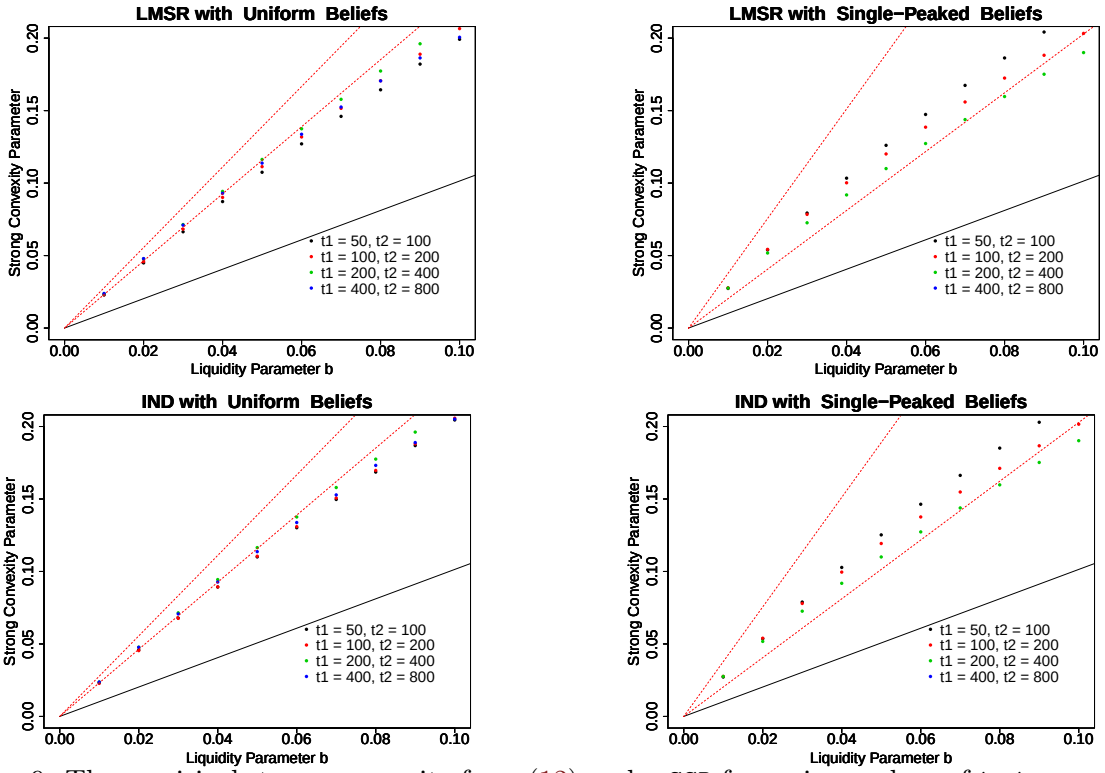

Figure 9: The empirical strong convexity from (12) under SSD for various values of t_1, t_2 represented as dots, asymptotic bounds for σ_{low} and σ_{high} from Theorem 7.6 (ignoring $O(b^2)$ terms) represented as lines. The red dashed lines are conjectured values for σ_{low} and σ_{high} given in Conjecture E.15.